2008-06

Internet topology generation based on reverse-engineered design principles performance tradeoffs between heuristic and optimization-based approaches

Derosier, Jonathan A.

Monterey, California. Naval Postgraduate School

http://hdl.handle.net/10945/4072

NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

INTERNET TOPOLOGY GENERATION BASED ON REVERSE-ENGINEERED DESIGN PRINCIPLES: PERFORMANCE TRADEOFFS BETWEEN HEURISTIC AND OPTIMIZATION-BASED APPROACHES

by

Jonathan A. Derosier

June 2008

Thesis Advisor:	David L. Alderson
Second Reader:	W. Matthew Carlyle

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			*Form Approved OMB No. 0704-0188*
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.			
1. AGENCY USE ONLY *(Leave blank)*	**2. REPORT DATE** June 2008	**3. REPORT TYPE AND DATES COVERED** Master's Thesis	
4. TITLE AND SUBTITLE Internet Topology Generation Based on Reverse-Engineered Design Principles: Performance Tradeoffs Between Heuristic and Optimization-Based Approaches			**5. FUNDING NUMBERS**
6. AUTHOR(S) Jonathan A. Derosier			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			**8. PERFORMING ORGANIZATION REPORT NUMBER**
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A			**10. SPONSORING/MONITORING AGENCY REPORT NUMBER**
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.			
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited			**12b. DISTRIBUTION CODE**

13. ABSTRACT (maximum 200 words)

The global Internet is a federation of computer networks that are owned and operated by Internet Service Providers (ISPs). Because ISPs do not share topology information for competitive and privacy reasons, researchers, operators, and policy makers who want to assess the performance and reliability of the system as a whole must infer structure from limited measurement data. We use reverse-engineering to infer underlying design principles of a national ISP and then develop models capable of generating ISP topologies ranging from regional to national scales. We contrast the behavior of optimal versus heuristic designs in terms of cost and performance. Unlike previous approaches that simply replicate observed network connectivity statistics, our approach yields networks that reflect the technological capabilities, economic constraints, operational requirements, and performance objectives faced by real ISPs. We complement our mathematics with computational tools that facilitate this network generation and analysis. To our knowledge, this thesis represents the first effort to incorporate these modeling principles in a process capable of generating realistic ISP networks at the national scale.

14. SUBJECT TERMS Internet, Topology Generation, Reverse Engineering, Optimization, IP, Network, Internet Service Provider, ISP, Motif, Heuristic, Optimal			**15. NUMBER OF PAGES** 91
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	**18. SECURITY CLASSIFICATION OF THIS PAGE** Unclassified	**19. SECURITY CLASSIFICATION OF ABSTRACT** Unclassified	**20. LIMITATION OF ABSTRACT** UU

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

INTERNET TOPOLOGY GENERATION BASED ON REVERSE-ENGINEERED DESIGN PRINCIPLES: PERFORMANCE TRADEOFFS BETWEEN HEURISTIC AND OPTIMIZATION-BASED APPROACHES

Jonathan A. Derosier
Captain, United States Marine Corps
B.S., Boston University, 2000

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

NAVAL POSTGRADUATE SCHOOL
June 2008

Author: Jonathan A. Derosier

Approved by: David L. Alderson
Thesis Advisor

W. Matthew Carlyle
Second Reader

James Eagle
Chairman, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

The global Internet is a federation of computer networks that are owned and operated by Internet Service Providers (ISPs). Because ISPs do not share topology information for competitive and privacy reasons, researchers, operators, and policy makers who want to assess the performance and reliability of the system as a whole must infer structure from limited measurement data. We use reverse-engineering to infer underlying design principles of a national ISP and then develop models capable of generating ISP topologies ranging from regional to national scales. We contrast the behavior of optimal versus heuristic designs in terms of cost and performance. Unlike previous approaches that simply replicate observed network connectivity statistics, our approach yields networks that reflect the technological capabilities, economic constraints, operational requirements, and performance objectives faced by real ISPs. We complement our mathematics with computational tools that facilitate this network generation and analysis. To our knowledge, this thesis represents the first effort to incorporate these modeling principles in a process capable of generating realistic ISP networks at the national scale.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I. INTRODUCTION 1
II. REVERSE ENGINEERING A NATIONAL ISP NETWORK 5
A. DATA 6
1. U.S. Census Bureau Data 6
2. Rocketfuel Data 6
B. ISP BACKBONE TOPOLOGY STRUCTURE 7
1. Routers and Links 7
2. Points of Presence (POP) 8
C. BACKBONE TOPOLOGY FOR AS 7018 10
D. POINT OF PRESENCE ROUTER STRUCTURE FOR AS 7018 14
1. Access Router Aggregation 14
2. Backbone Router Aggregation 15
3. Customer vs. Population 18
E. SUMMARY OF DESIGN PRINCIPLES 19
III. FORWARD ENGINEERING NETWORK TOPOLOGIES 21
A. PRE-PROCESSING: GATHERING NETWORK REQUIREMENTS 21
1. Customer and Access Router Assignment Model (CARAM) 22
B. BACKBONE TOPOLOGY GENERATION 24
1. Heuristic Backbone Router and Link Assignment Model 25
a. Backbone Router Assignment Model (BRAM) 25
b. Backbone Link Assignment Model (BLAM) 26
2. Optimal Backbone Topology Models 29
C. POST-PROCESSING: BUILDING A ROUTER-LEVEL MAP 38
IV. ANALYZING TOPOLOGIES 39
A. ANALYSIS DATA SETS 40
1. Subset 1: Small Network 42
2. Subset 2: Southern California Region 44
3. Subset 3: Three Large MSAs 46
4. Subset 4: Western United States 48
5. Subset 5: North Eastern United States 50
6. Subset 6: United States Edge Heavy 52
7. Subset 7: United States Core Heavy 54
8. Subset 8: All MSAs 56
B. ANALYSIS RESULTS 58
V. CONCLUSIONS 65
APPENDIX 67
LIST OF REFERENCES 69
INITIAL DISTRIBUTION LIST 71

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF FIGURES

Figure 1. The structure of an ISP Network Topology reflects the functional need to support its customers. 3
Figure 2. An Autonomous System (AS) is an IP network under single administrative control. The Internet is a collection of interconnected ASes. Connections between ASes represent peering relationships. 5
Figure 3. Conceptual Representation of an ISP Backbone Topology. 9
Figure 4. Backbone Topology for AS 7018 Rocketfuel Data. Nodes represent Points of Presence (POP). The Core POPs are labeled with their DNS location code. Links represent at least one logical connection between a pair of routers, each in a different POP. The topology reflects a hub and spoke design motif. 11
Figure 5. Distribution of Customer Connections per Access Router for AS 7018. 15
Figure 6. Router connectivity within an individual POP. (a) Two-Backbone Router POP. (b) Four-Backbone Router POP. (c) Six-Backbone Router POP. 16
Figure 7. Number of Gigabit Backbone Routers vs. Number of Access Routers for each Point of Presence in AS 7018. The number in each data point is the number of POPs observed with that combination of access routers and backbone routers. 17
Figure 8. Network Topology Generation Process. 21
Figure 9. Pre-Processing Stage. This stage is applied to each MSA in turn. Inputs are the population and market penetration of each MSA. The outputs are the number of customers and access routers at each node. 22
Figure 10. The Customer and Access Router Assignment Model (CARAM) prescribes the number of customers and access routers at a node given the population and market penetration at the node. The number of customers at a node is linearly proportional to the weighted population. The number of access routers assigned is dependent upon the number of customers. (i.e., a node with 500 customers would have 13 access routers with 38 customers per access router) 23
Figure 11. Backbone Topology Generation Inputs and Outputs. 24
Figure 12. Backbone Router Assignment Model. $g_1 = 4$, $g_2 = 7$, $g_3 = 12$. 25
Figure 13. Link Elimination Procedure. We consider each combination of three nodes. The links between the nodes form the legs of a triangle. If the length of the longest leg of the triangle, multiplied by the parameter $\alpha \in [1.0, 2.0]$, is greater than the sum of the lengths of the two shortest legs, then eliminate the link associated with the longest leg. 27
Figure 14. Networks generated with different values of λ 28

Figure 15. First layer core networks generated with different values of α 28
Figure 16. Cost and Capacity Assumption for each Node Type. The internal structure for each node follows directly from the design motifs observed in AS 7018 and illustrated in Figure 6. 30
Figure 17. Post Processing Stage 38
Figure 18. We build a router-level topology from the backbone topology in the post processing stage. a. Backbone representation with (2) two-backbone router core nodes. b. Equivalent router-level topology. 38
Figure 19. Analysis Methodology... 39
Figure 20. Metropolitan Statistical Area Subsets 1-8... 41
Figure 21. Subset 1 Backbone Topology Generation Solutions. a. Heuristic Model Solution. b. Optimal Maximum Flow Modes (Equal Cost) solution. c. Optimal Minimum Cost Model (Equal Flow) solution. 43
Figure 22. Subset 2 Backbone Topology Generation Solutions. a. Heuristic Model Solution. b. Optimal Maximum Flow Modes (Equal Cost) solution. c. Optimal Minimum Cost Model (Equal Flow) solution. 45
Figure 23. Subset 3 Backbone Topology Generation Solutions. a. Heuristic Model Solution. b. Optimal Maximum Flow Modes (Equal Cost) solution. c. Optimal Minimum Cost Model (Equal Flow) solution. 47
Figure 24. Subset 4 Backbone Topology Generation Solutions. a. Heuristic Model Solution. b. Optimal Maximum Flow Modes (Equal Cost) solution. c. Optimal Minimum Cost Model (Equal Flow) solution. 49
Figure 25. Subset 5 Backbone Topology Generation Solutions. a. Heuristic Model Solution. b. Optimal Maximum Flow Modes (Equal Cost) solution. c. Optimal Minimum Cost Model (Equal Flow) solution. 51
Figure 26. Subset 6 Backbone Topology Generation Solutions. a. Heuristic Model Solution. b. Optimal Maximum Flow Modes (Equal Cost) solution. c. Optimal Minimum Cost Model (Equal Flow) solution. 53
Figure 27. Subset 7 Backbone Topology Generation Solutions. a. Heuristic Model Solution. b. Optimal Maximum Flow Modes (Equal Cost) solution. c. Optimal Minimum Cost Model (Equal Flow) solution. 55
Figure 28. Subset 8 Backbone Topology Generation Solutions. a. Heuristic Model Solution. b. Optimal Maximum Flow Modes (Equal Cost) solution. c. Optimal Minimum Cost Model (Equal Flow) solution. 57
Figure 29. Router-level topology cost and flow comparisons for subsets 1-8. a. Cost comparison. b. Total throughput comparison using successive shortest path routing. 59
Figure 30. Achieved Customer Bandwidth (Shortest Path Routing) 61
Figure 31. Achieved Customer Bandwidth (Successive Shortest Path Routing) . 61
Figure 32. Individual Backbone Router Utilization Under Maxflow Conditions. The backbone routers include routers from all eight MSA subsets and each backbone generation model. Backbone router utilization depends upon the topology structure and traffic engineering used in the network. The wide variation in utilization with a majority of routers being used indicates reasonable resource allocation.

Nearly all (99.65%) of backbone routers are utilized, some more then others, with the vast majority under 50% utilization. 62
Figure 33. Individual Access Router Utilization Under Maxflow Conditions. Because access routers demand traffic in proportion to the number of customers, the total utilization of access routers increases linearly with customer count. Routers from each subset and backbone model (heuristic, equal cost or equal flow) lie on separate lines. 63

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

Table 1. Router counts for AS 7018 Point of Presences with corresponding Metropolitan Statistical Area population. 12
Table 2. Observed features in the AS 7018 backbone topology and their engineering design reasoning. 19
Table 3. Observed features in the AS 7018 point of presence structure and their engineering design reasoning.. 20
Table 4. Optimal Backbone Topology Model Arc Types.................................. 31
Table 5. Metropolitan Statistical Area Subset Summary 40
Table 6. Model Cost and Capacity Parameters.. 42
Table 7. Subset 1 Results 42
Table 8. Subset 2 Results 44
Table 9. Subset 3 Results 46
Table 10. Subset 4 Results 48
Table 11. Subset 5 Results 50
Table 12. Subset 6 Results 52
Table 13. Subset 7 Results 54
Table 14. Subset 8 Results 56
Table 15. Throughput achieved by the router topology representation relative to the backbone topology representation.. 58
Table 16. Router and Arc Counts for Router-Level Topologies.......................... 59

THIS PAGE INTENTIONALLY LEFT BLANK

EXECUTIVE SUMMARY

The Internet is a critical component of our economic and social fabric, and many civilian and military systems are dependent upon it in one way or another. The foundation of the Internet is the physical network of computers, routers, and fiber optical lines connecting them. Internet Service Providers (ISPs), the owners and operators of these networks, do not publish their topology information, and thus researchers, IT professionals, and even ISP operators do not know the Internet's large-scale topology structure. To fill this void, researchers use experimental methods to measure and infer the router-level structure of the Internet.

One popular approach to characterizing router-level network structure is to apply graph theoretic and/or statistical techniques to the connectivity patterns observed in measurement experiments. These characterizations are typically accompanied by generative models that faithfully reproduce the observed statistics. This approach leads to descriptive models of network structure that, while interesting, typically fail to reveal explanatory or causal relationships at work in the design and operation of real ISP networks.

This thesis follows an alternative approach in which the causal forces shaping network design and deployment are reflected in an optimization problem. This type of *optimization-based reverse engineering* has roots in previous work, but this thesis represents the first effort to incorporate these modeling principles in a process capable of representing a router-level network at a national scale.

Using this alternative modeling approach, we seek to design router-level topologies that provide sufficient and reliable bandwidth to network customers at a reasonable cost. To accomplish this, we do three things. One, we analyze an existing router-level topology for a U.S. National Tier-1 ISP and *reverse engineer* its key design principles (e.g., backbone routers occurring in pairs for redundancy). Two, we *forward engineer* a network topology generation process

based upon the design principles that we observe. In this generation process, we develop both heuristic and optimal generation methods. Finally, we validate that the network topologies provide sufficient bandwidth and are realistic based on what we currently know about network topologies. In addition, we compare and contrast heuristic and optimally generated topologies to quantify their differences in terms of cost and performance.

We generate networks for eight different customer populations that range from small regional populations, e.g., Southern California, to the National level, e.g., the entire United States. For each customer population we generate three topologies, one using the heuristic method, one using an optimization model that maximizes throughput subject to a budget, and a third using an optimization model that minimizes cost subject to a throughput requirement.

We compare the network topologies based on two measures of performance: cost, and throughput. Cost is sum of the cost of each network component (routers and links) in the router-level topology and is measured in thousands of dollars ($K). Throughput is represented by the sum of the flow across all pairs of communicating routers on the network and is measured by *bandwidth* in gigabits per second (Gps). To represent fair traffic demand we assume a *gravity* model, which constrains the demand between each pair of communicating routers to be proportional to the product of their customer populations.

There are three main contributions of this thesis. First, it presents a systematic process by which one can generate a “realistic, yet fictitious” ISP networks at a national scale. The topologies generated from our process are realistic, in the sense that (1) they adhere to basic technological and economic constraints facing the design of real ISP networks; (2) they are derived from real geographic and population data representing real customer markets; and (3) they are generated at the level of individual routers, meaning that these networks can be used as a basis for packet-level simulations of Internet traffic.

The second main contribution of this thesis is the quantitative comparison of heuristic and optimal topology generation schemes, in terms of network performance and cost. We use these results to develop insight into the tradeoffs between optimal and heuristic design philosophies at work in real ISPs.

Third, we support our analytic and numerical results with an automated decision support tool developed in Excel/VBA and using state-of-the-art commercial optimization software (GAMS/CPLEX). This integrated tool allows its user to conveniently select customer markets at the national scale, design and illustrate a high-level “backbone” ISP network, and then generate the corresponding router-level topology. To date, comparable topology generation tools do not exist within the scientific community.

Collectively, this thesis provides researchers and operators with the mathematical framework and computational tools necessary to explore the relationship between ISP structure and function, both for the current operational environment and in the future.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

The Internet is a critical component of our economic and social fabric and many civilian and military systems are dependent upon it in one way or another. The global Internet is a federation of independently owned and operated computer networks that support a standard suite of communication protocols. Internet Service Providers (ISPs) are the owner-operators of these networks. ISPs are classified into tiers based on *peering* (settlement free interconnection) relationships. Tier-1 ISPs peer with every other Tier-1 ISP and therefore can reach any network on the Internet without purchasing transit. AT&T and Sprint are examples of U.S. National Tier-1 ISPs. Entities within the Department of Defense are also ISPs in the sense that they build and operate a variety of global networks running the Internet protocol suite and are connected to other ISP networks.

The foundation of the Internet is the physical network of computers, routers, and fiber-optical lines connecting them. The design of this *router-level network* is important because it directly affects the overall cost, reliability, and performance of the system. The connectivity within a router-level network is not arbitrary or random; rather, it follows from design that has specific structure to support communication between the network's customers. The relationship between the customer population and the network topology reflects many elements such as technological capabilities, economic constraints, performance objectives, and any design methodologies in use.

Over the past decade, there has been considerable interest in understanding the large-scale structure of the Internet at the router-level and at other levels of abstraction. Because ISPs regard their network topologies as a source of competitive advantage, they are reluctant to share topology information, thereby leaving researchers, IT professionals, and even ISP operators in the dark about the structure of the router-level Internet as a whole.

To overcome the lack of publicly available Internet topology data, researchers have developed a variety of techniques to infer network structure from measurement experiments. These techniques use well-understood software tools, such as `traceroute`, to measure traffic as it traverses the network. This measurement data is then analyzed with the hope of identifying key structural features that dictate network performance, robustness, and vulnerability.

One popular approach to characterizing router-level network structure has been to apply graph theoretic and/or statistical techniques to the connectivity patterns observed in measurement experiments. These characterizations are typically accompanied by generative models that faithfully reproduce the observed statistics (Li et al. 2004). While this approach leads to descriptive models of network structure that are interesting and provocative, it typically fails to reveal explanatory or causal relationships at work in the design and operation of real ISP networks. Owing to the inherent diversity among networks sharing the same statistics, the ability of a single model to replicate observed statistics provides little validation that it is accurate or even realistic (Alderson, 2008).

This thesis follows an alternative approach in which the causal forces shaping network design and deployment are reflected in an optimization problem. The roots of this type of *optimization-based reverse engineering* can be traced to Alderson et al. (2003) and Alderson et al. (2005), but this thesis represents the first effort to incorporate these modeling principles in a process capable of representing a router-level network at a national scale.

A fundamental challenge with this alternative approach is that network design problems are inherently hard to solve optimally, and so heuristics are often used in practice. However, it is unclear what potential cost is being paid by using heuristic solutions. In other words, what tradeoffs in performance and cost exist between optimally and heuristically designed networks?

This thesis explores the relationship between customer population and network topology in two ways. First, in Chapter II, we study the topology of a real Tier-1 ISP and, using census data, we infer the way in which design patterns, or *motifs*, support functional needs in terms of throughput and reliability. We refer to this process as *reverse engineering*. Then, in Chapter III, we use the inferred relationships as the basis for a *forward engineering* design process that generates optimal network topologies under competing objectives of performance and cost. In Chapter IV, we compare the output from these two approaches for eight different case studies, ranging from U.S. regional to national markets. We summarize our results and describe opportunities for future work in Chapter V.

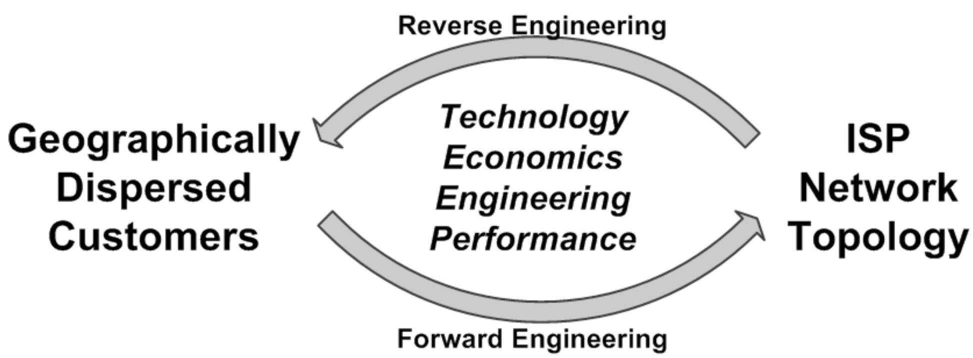

Figure 1. The structure of an ISP Network Topology reflects the functional need to support its customers.

There are three main contributions of this thesis. First, this thesis presents a systematic process by which one can generate a "realistic, yet fictitious" ISP networks at a national scale. The topologies generated from our process are realistic, in the sense that (1) they adhere to basic technological and economic constraints facing the design of real ISP networks; (2) they are derived from real geographic and population data representing real customer markets; and (3) they are generated at the level of individual routers, meaning that these networks can be used as a basis for packet-level simulations of Internet traffic.

The resulting network topologies are dramatically different in structure and fidelity then currently popular topology generation schemes that replicate statistical network features (Li et al., 2004).

The second main contribution of this thesis is the quantitative comparison of heuristic and optimal topology generation schemes, in terms of network performance and cost. We use these results to develop insight into the tradeoffs between optimal and heuristic design philosophies at work in real ISPs.

Third, we support our analytic and numerical results with an automated decision support tool developed in MS Excel with Visual Basic for Applications (VBA) and using state-of-the-art commercial optimization software (GAMS/CPLEX). This integrated tool allows its user to conveniently select customer markets at the national scale, design and illustrate a high-level “backbone” ISP network, and then generate the corresponding router-level topology. To date, comparable topology generation tools do not exist within the scientific community.

Collectively, this thesis provides researchers and operators with the mathematical framework and computational tools necessary to explore the relationship between ISP structure and function, both for the current operational environment and in the future.

II. REVERSE ENGINEERING A NATIONAL ISP NETWORK

Our approach to router-level topology modeling begins with three assumptions. First, we assume that a network topology is not random but has structural features that support the functional requirements of the network's customer population. Second, we assume that the structure of the topology reflects heuristic design patterns, or *motifs*, used by the engineers of the network to design it. Third, we assume that these design motifs can be inferred using an existing network topology and its supported population.

An Autonomous Systems (AS) is an IP network under single administrative control. That is, an AS has a single decision maker (administrator) who is responsible for the provisioning, traffic engineering, and routing policies that are seen by the rest of the Internet. We focus our research on the AS because it as this level of abstraction that network topology design decisions are made. Although a Tier-1 ISP may have one or more ASes, we will use the terms AS and ISP interchangeably in this thesis. We illustrate the Internet as a collection of interconnected ASes in Figure 2. In this chapter, we infer the design motifs for AS 7018, a national network owned and operated by AT&T.

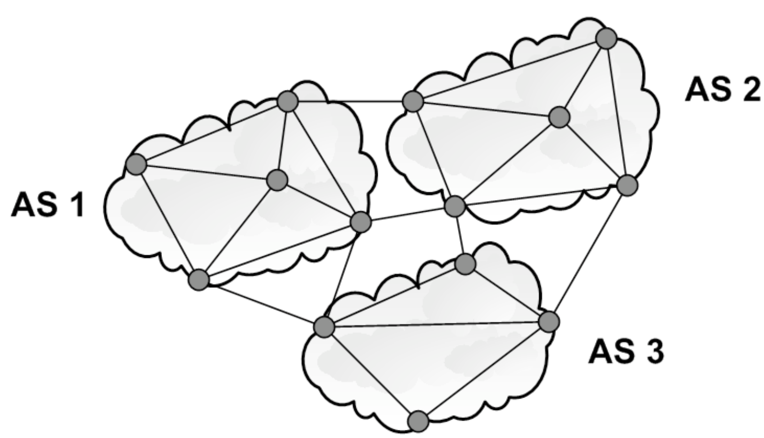

Figure 2. An Autonomous System (AS) is an IP network under single administrative control. The Internet is a collection of interconnected ASes. Connections between ASes represent peering relationships.

A. DATA

1. U.S. Census Bureau Data

The United States Census Bureau maintains population data categorized by geographic subdivisions. Cities, Counties, and Metropolitan Statistical Areas are three principle subdivisions.

A Metropolitan Statistical Area (MSA) is a central urbanized area—a contiguous area of relatively high population density. An MSA consists of a collection of counties that are connected by strong social and economic ties as measured by commuting and employment (U.S. Census Bureau, 2007).

We use MSAs to represent regional markets for ISPs.

2. Rocketfuel Data

We derive design motifs from router-level topology data for AS 7018 as it was collected circa 2003. The data is publicly available and was collected by the Rocketfuel Project (Spring et al., 2003), an ISP topology-mapping tool that uses focused `traceroute` experiments to infer the internal router-level structure of a single ISP. The Rocketfuel project has mapped several ISPs within the United States, Europe and Australia. For each AS studied, Rocketfuel data includes information about routers (type, geographic location, etc.) and the links connecting them. Although the Rocketfuel maps are not 100% accurate, they have been broadly validated and are considered among the best of currently available router-level topology maps.

B. ISP BACKBONE TOPOLOGY STRUCTURE

1. Routers and Links

Routers are the building blocks of computer networks. Routers are specialized computers that receive incoming network traffic and forward it appropriately to its next destination. Routers are connected by physical wires (e.g., optical fibers or copper wires). For long-haul traffic, a network of optical fibers comprises the optical layer of the network upon which higher layers of the network are built. From an Internet Protocol (IP) perspective, routers are connected by logical links. An *IP link* represents one-hop IP connectivity between two routers. Throughout this thesis, all links are IP links.

Routers vary widely based on their purpose, but for a Tier-1 ISP they can be broadly categorized as either *backbone* or *access* routers. *Backbone* routers exist within an AS and communicate primarily to routers belonging to the AS. They typically support few high bandwidth links and serve to interconnect backbone routers over long distances, or as aggregation points for access routers. *Access* routers communicate internally to an AS’s core routers and externally to customers. They typically support many low bandwidth links on the customer side and connect to a few backbone routers in the network's backbone.

2. Points of Presence (POP)

A *Point of Presence* (POP) is a collocated logical collection of routers that serves primarily as an access point for the network's customers. The POPs in an AS are geographically distributed and each correspond roughly to a regional market. Every router in an AS belongs to a POP. Access routers within a POP serve as the physical connection between the ISP and its customers.

Some POPs have backbone routers in addition to access routers. This infrastructure can be thought of as "sitting atop" the access infrastructure. The backbone routers within select POPs are interconnected by high capacity links that span relatively large distances to other POPs.

Every access router within a POP must connect to a backbone router. When backbone routers are collocated with access routers, this connection is internal to the POP. In POPs that do not have a backbone router, the access routers must connect to a backbone router in a nearby POP.

Throughout this thesis, we use the following terminology when referring to the backbone topology.

- A *Core POP* is a POP that has backbone routers.
- An *Edge POP* is a POP that does not have backbone routers.
- A *Link* is one or more logical connections between two routers, each in a different POP.
- An *Access-Backbone Link* is a link between an Edge POP and a Core POP.
- A *Backbone-Backbone Link* is a link between two Core POPs.

We illustrate a backbone topology structure in Figure 3.

Points of Presence (POPs) represent the geographic locations where an ISP connects to its customers. POPs contain access routers---the physical connection devices. We illustrate POPs as light gray spheres.	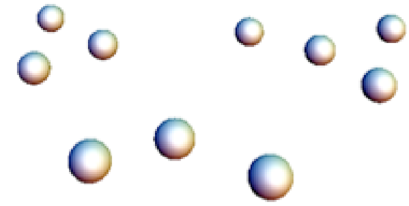
The backbone of an ISP's network is built from additional routing infrastructure located in select POPs. We refer to these as *Core POPs*. POPs with only access routers are *Edge POPs*.	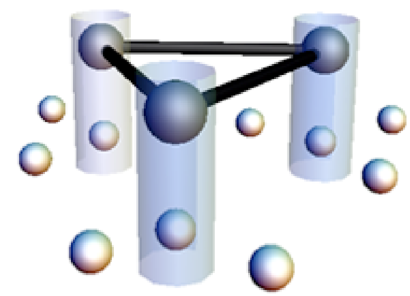
All access routers in the POPs connect to the backbone either internally (Core POPs) or externally (Edge POPs).	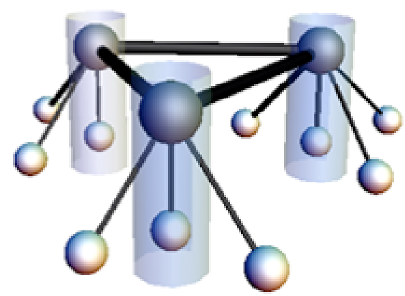
Viewed from above Core POPs appear as "hubs" and edge POPs appear as "spokes". We refer to this structure as *hub and spoke*.	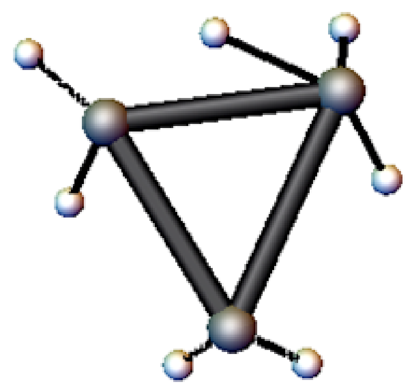

Figure 3. Conceptual Representation of an ISP Backbone Topology.

C. BACKBONE TOPOLOGY FOR AS 7018

We illustrate the backbone topology for AS 7018 as measured by Rocketfuel in Figure 4. Dark nodes represent the Core POPs and light nodes represent the Edge POPs.

The POPs in AS 7018 correspond reasonably well to MSAs. Larger MSAs may have multiple POPs in them. In these cases, only one of these POPs has backbone routers and the vast majority of the access routers. An example of this is Chicago, which has POPs `cgcil`, `chcil`, `chgil`, and `okbil` with (16, 1, 1, 1) access routers and (6, 0, 0, 0) backbone routers, respectively. The population and router counts for AS 7018's POPs and corresponding MSAs, sorted by population, are listed in Table 1.

Edge POPs have an average of 1.1 connections indicating that edge POPs typically connect only to a single core POP. Core POPs support an average of 5.3 edge POPs and connect to an average of 3.8 core POPs. This structure is characteristic of a "hub and spoke" design motif.

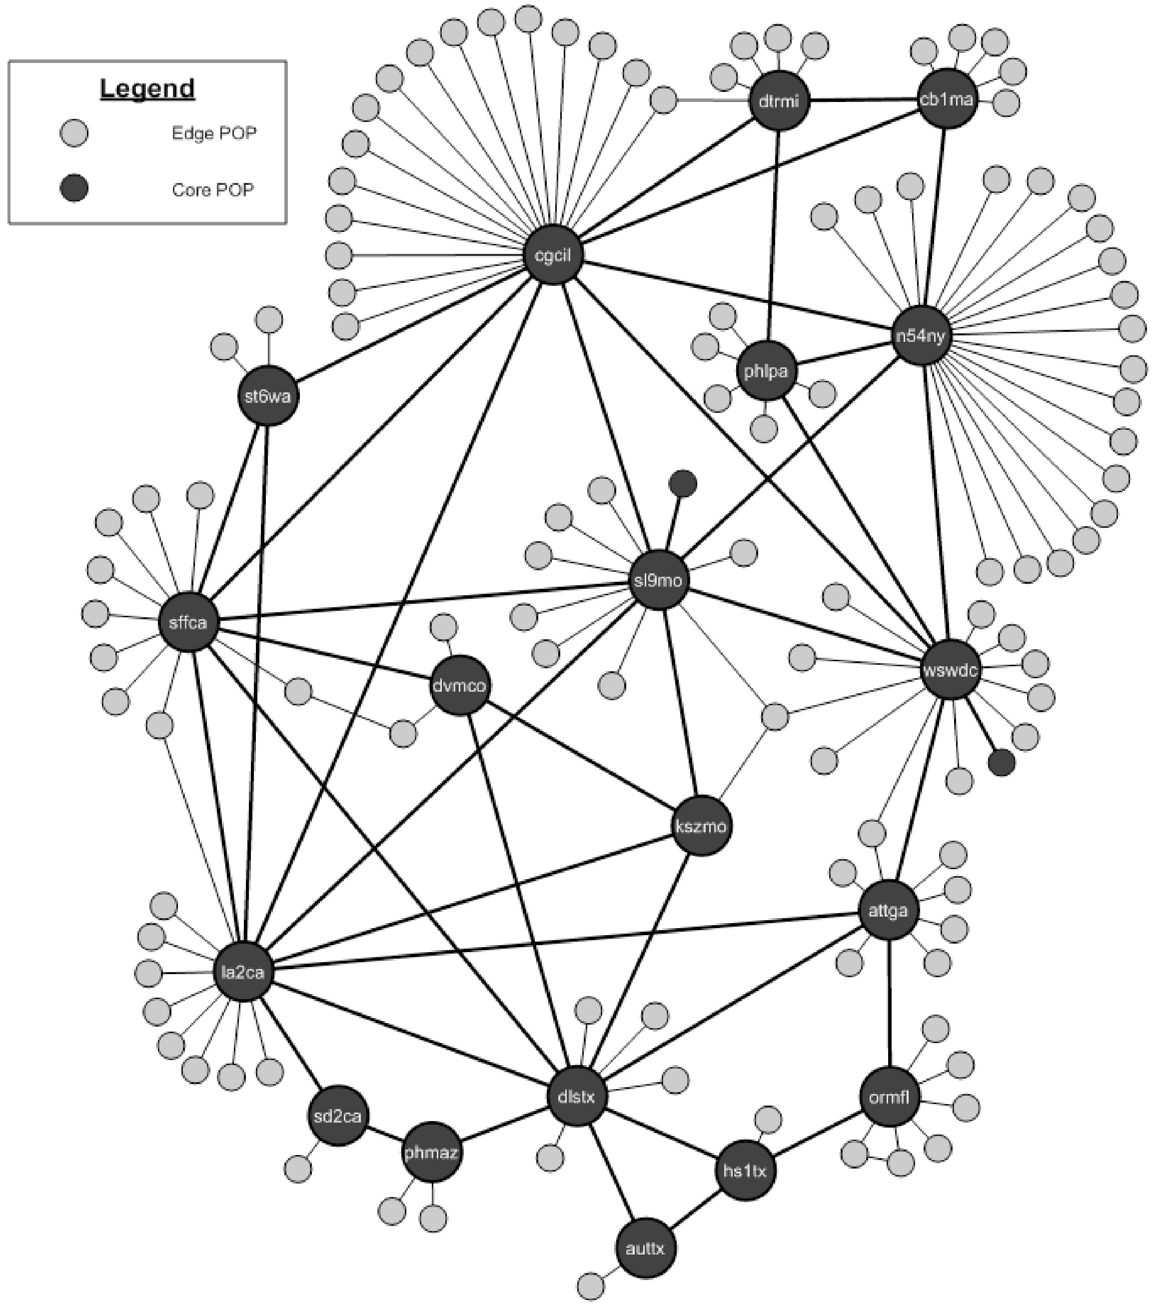

Figure 4. Backbone Topology for AS 7018 Rocketfuel Data. Nodes represent Points of Presence (POP). The Core POPs are labeled with their DNS location code. Links represent at least one logical connection between a pair of routers, each in a different POP. The topology reflects a hub and spoke design motif.

Table 1. Router counts for AS 7018 Point of Presences with corresponding Metropolitan Statistical Area population.

Metropolitan Statistical Area	**gbr**	**ar**	**Population [2000 Census]**
New York, NY	6	26	11,296,377
Los Angeles, CA	6	15	9,519,338
Chicago, IL	6	19	7,628,412
Houston-Sugar Land-Baytown, TX	2	4	4,715,407
Atlanta-Sandy Springs-Marietta, GA	6	13	4,247,981
Philadelphia, PA	2	4	3,849,647
Washington D.C.	6	13	3,727,565
Dallas, TX	6	13	3,451,226
Riverside-San Bernardino-Ontario, CA	0	3	3,254,821
Phoenix-Mesa-Scottsdale, AZ	2	5	3,251,876
Minneapolis-St. Paul-Bloomington, MN	0	3	2,968,806
Anaheim, CA	0	2	2,846,289
San Diego-Carlsbad-San Marcos, CA	2	5	2,813,833
Long Island, NY	0	1	2,753,913
St. Louis, MO	8	11	2,721,491
Baltimore-Towson, MD	0	1	2,552,994
Pittsburgh, PA	0	2	2,431,087
Tampa-St. Petersburg-Clearwater, FL	0	3	2,395,997
Oakland, CA	0	1	2,392,557
Warren, MI	0	1	2,391,395
Seattle, WA	4	7	2,343,058
Mimai, FL	0	4	2,253,362
Edison, NJ	0	2	2,173,869
Denver-Aurora, CO	4	8	2,157,756
Cleveland-Elyria-Mentor, OH	0	3	2,148,143
Newark, NJ	0	4	2,098,843
Detroit, MI	2	4	2,061,162
Cincinnati-Middletown, OH	0	1	2,009,632
Portland-Vancouver-Beaverton, OR	0	2	1,927,881
Kansas City, MO	2	2	1,836,038
Boston, MA	4	9	1,812,937
San Jose-Sunnyvale-Santa Clara, CA	0	3	1,735,819
San Francisco, CA	6	15	1,731,183
San Antonio, TX	0	1	1,711,703
Fortworth, TX	0	2	1,710,318
Orlando-Kissimmee, FL	4	11	1,644,561
Fort Lauderdale, FL	0	3	1,623,018
Providence-New Bedford-Fall River, RI	0	1	1,582,997
Virginia Beach-Norfolk-Newport News, VA	0	2	1,576,370
Indianapolis-Carmel, IN	0	2	1,525,104

Metropolitan Statistical Area	gbr	ar	Population [2000 Census]
Milwaukee-Waukesha-West Allis, WI	0	2	1,500,741
Cambridge, MA	0	1	1,465,396
Las Vegas-Paradise, NV	0	1	1,375,765
Charlotte-Gastonia-Concord, NC	0	2	1,330,448
New Orleans-Metairie-Kenner, LA	0	2	1,316,510
Nashville-Davidson, TN	0	2	1,311,789
Austin-Round Rock, TX	2	3	1,249,763
Memphis, TN	0	1	1,205,204
Camden, NJ	0	1	1,186,999
Buffalo-Niagara Falls, NY	0	1	1,170,111
Louisville/Jefferson County, KY	0	1	1,161,975
Hartford-West Hartford-East Hartford, CT	0	2	1,148,618
West Palm Beach, FL	0	1	1,131,184
Jacksonville, FL	0	1	1,122,750
Richmond, VA	0	1	1,096,957
Oklahoma City, OK	0	2	1,095,421
Bethesda, MD	0	1	1,068,618
Birmingham-Hoover, AL	0	1	1,052,238
Rochester, NY	0	1	1,037,831
Salt Lake City, UT	0	2	968,858
Bridgeport-Stamford-Norwalk, CT	0	2	882,567
Honolulu, HI	0	1	876,156
Tulsa, OK	0	1	859,532
Dayton, OH	0	1	848,153
Tucson, AZ	0	1	843,746
Albany-Schenectady-Troy, NY	0	2	825,875
Raleigh-Cary, NC	0	2	797,071
Omaha-Council Bluffs, NE	0	2	767,041
Worcester, MA	0	1	750,963
Grand Rapids-Wyoming, MI	0	1	740,482
Albuquerque, NM	0	1	729,649
Akron, OH	0	1	694,960
Syracuse, NY	0	1	650,154
Columbia, SC	0	1	647,158
Greensboro-High Point, NC	0	2	643,430
Little Rock-North Little Rock-Conway, AR	0	1	610,518
Colorado Springs, CO	0	2	537,484
Harrisburg-Carlisle, PA	0	2	509,074
Madison, WI	0	1	501,774
Portland-South Portland-Biddeford, ME	0	1	487,568
Des Moines-West Des Moines, IA	0	1	481,394
Spokane, WA	0	1	417,939
Manchester-Nashua, NH	0	1	380,841
Davenport-Moline-Rock Island, IA	0	3	376,019
Springfield, MO	0	1	368,374

Metropolitan Statistical Area	gbr	ar	Population [2000 Census]
Trenton-Ewing, NJ	0	2	350,761
South Bend-Mishawaka, IN	0	1	316,663
Lynchburg, VA	2	0	228,616
Champaign-Urbana, IL	0	2	210,275

D. POINT OF PRESENCE ROUTER STRUCTURE FOR AS 7018

A POP is designed to aggregate the traffic from many low bandwidth customer links into a few high bandwidth inter-POP links. This aggregation occurs at the access and backbone routers. The interconnection of routers within a single POP, reflects a redundant hierarchal design motif.

1. Access Router Aggregation

Access routers aggregate traffic between customer routers and backbone routers. In AS 7018, access routers have two parallel upstream connections, one each to a backbone router, providing for upstream redundancy, and some number of downstream customer router connections. The distribution of downstream customer router connections per access router is shown in Figure 5. The distribution reinforces that access routers can support a finite number of customer connections. This distribution is uni-modal and reasonably symmetric. The lower and upper quartiles occur at 20 and 40 customer connections. Engineering can explain the tails of the distribution. The lower tail may represent incomplete data, where not all connections on a router are observed, routers that support very few customers perhaps in remote sites with few customers, or new routers that have not been fully loaded. The upper tail may represent routers that are overloaded perhaps to defer the cost of installing additional routers.

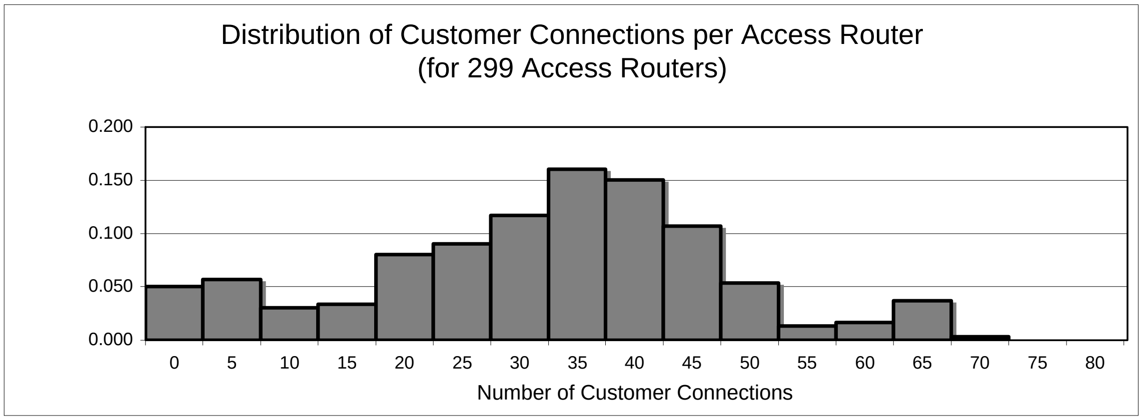

Figure 5. Distribution of Customer Connections per Access Router for AS 7018.

2. Backbone Router Aggregation

Backbone routers aggregate traffic from access routers into a few high bandwidth inter-POP connections. For AS 7018, we observe that if backbone routers are present within a POP, they occur in pairs, and the backbone router configuration reflects the number of backbone routers in the POP (two, four or six). These configurations are illustrated in Figure 6. The backbone routers need to support both downstream access router connections and upstream inter-POP backbone router connections. A two-backbone router configuration supports both downstream and upstream connections from the same pair of backbone routers. Different pairs of backbone routers handle the downstream and upstream connections in four-backbone and six-backbone router configurations.

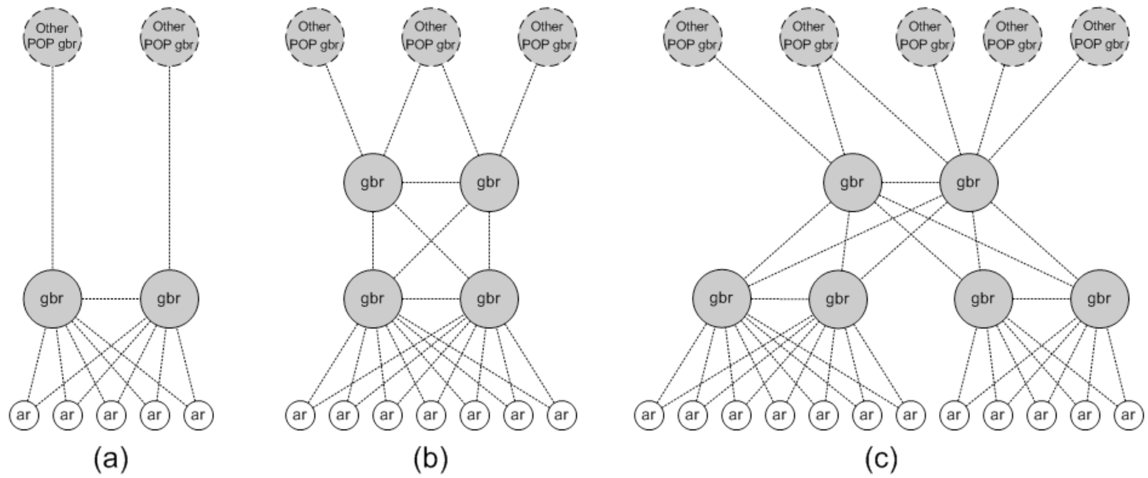

Figure 6. Router connectivity within an individual POP. (a) Two-Backbone Router POP. (b) Four-Backbone Router POP. (c) Six-Backbone Router POP.

In AS 7018, the number of backbone routers is closely related to the number of access routers, increasing as the number of access routers increases as illustrated in Figure 7. The edge POPs always have less then four access routers. All but three core POPs have three or more access routers. The three exceptions are known legacy sites, supporting dial-up access and other types of connectivity.

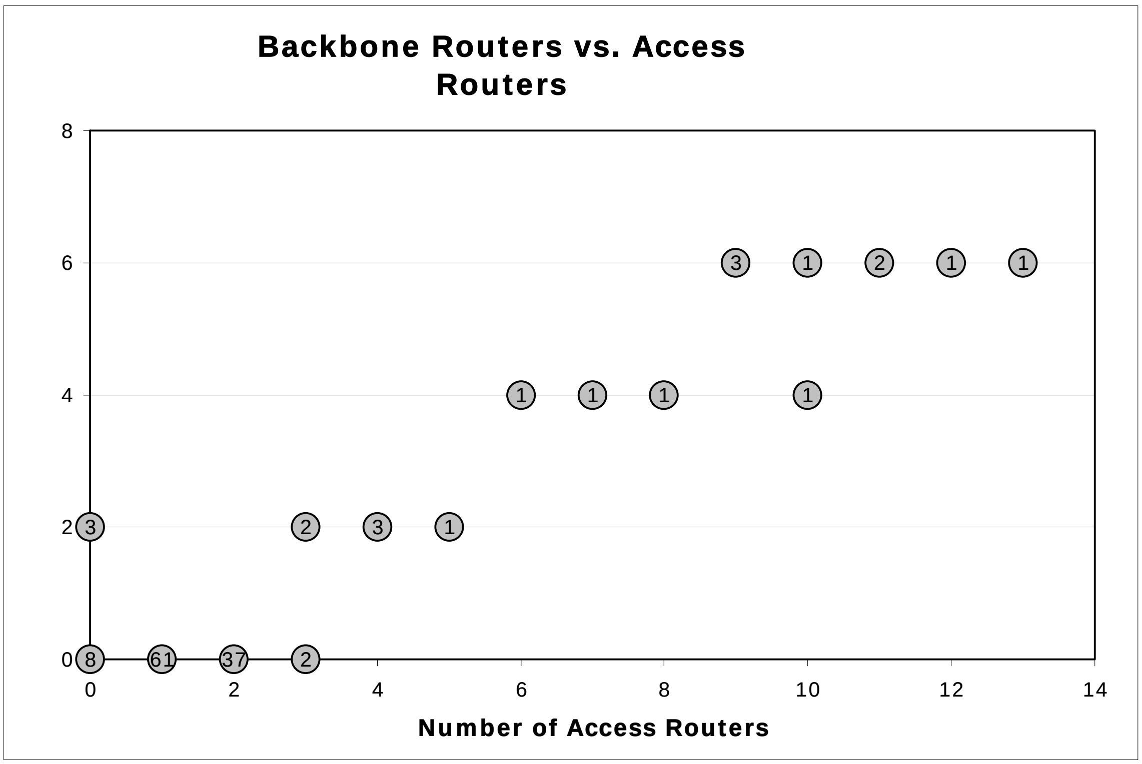

Figure 7. Number of Gigabit Backbone Routers vs. Number of Access Routers for each Point of Presence in AS 7018. The number in each data point is the number of POPs observed with that combination of access routers and backbone routers.

3. Customer vs. Population

The number of customer connections in an ISP's POP reflects local population and market penetration. POPs in locations with high populations tend to have more customers, access routers, and backbone routers. However, for AS 7018, we observed no linear relationship between population and the number of backbone routers, number of access routers, or number of customers. We assume then that the ISP has a different *market penetration* for each MSA that relates the number of network customers in the MSA to the census population of the MSA.

E. SUMMARY OF DESIGN PRINCIPLES

We conclude this chapter with a list of the structural features that we observe in Rocketfuel data for AS 7018. These features make clear sense in the context of engineering design and so we use them as design principles in our forward engineering process.

Table 2. Observed features in the AS 7018 backbone topology and their engineering design reasoning.

Observed Feature	**Engineering Design Reasoning**
-POPs can be divided into two distinct classes: those with backbone routers and those without backbone routers.	While all POPs aggregate traffic, only some POPs support backbone infrastructure (Core).
-POPs without backbone routers typically have one POP-POP link. This link is to the nearest POP that has backbone routers.	It is more efficient to connect the access routers in a small POP to the backbone routers in a nearby larger POP then to build and maintain backbone structure at a small POP.
-POPs with backbone routers typically have many POP-POP links. These links connect to POPs that have no backbone routers and to POPs that have backbone routers.	Backbone POPs serve as hubs in “hub and spoke” design motif.

Table 3. Observed features in the AS 7018 point of presence structure and their engineering design reasoning.

Observed Feature	**Engineering Design Reasoning**
-A POP can have zero, two, four, or six backbone routers.	Backbone routers occur in pairs for redundancy.
-The POP structure is related to the number of backbone routers in the POP.	The backbone router configuration within a POP determines its bandwidth capacity.
-The number of backbone routers is related to the number of access routers in the POP.	The backbone routers serve to aggregate traffic from the access routers. Therefore the number of access routers drives the backbone router requirements.
-The POP structure is scalable, i.e., the two-backbone router structure is contained within the four-backbone router structure and the four-backbone router structure is contained within the six-backbone router structure.	Scalable structure supports the expansion of POPs as more capacity is required.
-An access router connects in parallel to a pair of backbone routers.	Connecting in parallel provides redundancy in case of a backbone router or link failing
-An access router can support a finite number of customers.	Router degree is constrained by the number of line cards it can support. Line cards have a port/bandwidth configuration.

III. FORWARD ENGINEERING NETWORK TOPOLOGIES

In this chapter, we develop a process for generating ISP network topologies using the structural features observed in the AS 7018 network as a template. We start by grouping customer populations by geographical regions. Our objective is then to construct a network topology that provides reliable and sufficient connectivity for the ISP's customer population at a reasonable cost. The generation process is comprised of the three sequential stages illustrated in Figure 8.

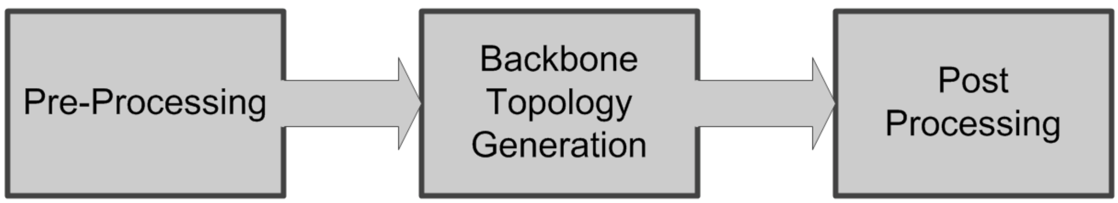

Figure 8. Network Topology Generation Process.

Backbone topology generation is the central focus of this thesis. The design of the backbone topology fundamentally impacts the cost, throughput and robustness of the network. We develop both heuristic and optimal methods for designing the backbone topology. We also develop pre-processing and post-processing stages to infer parameters and work with real data. We apply the same pre-processing and post-processing to all networks.

A. PRE-PROCESSING: GATHERING NETWORK REQUIREMENTS

In the Pre-Processing Stage, we associate a node with each geographical region (MSA), identify the customer demand for the MSA, and choose the access router interconnection structure to support that demand. Inputs to the pre-processing stage include the census population and an assumed market penetration for each MSA in the network. Outputs of the pre-processing stage include the assumed number of customers and number of access routers at each node. We define demand as the number of customers per access router. We

generate the number of customers and access routers using the Customer and Access Router Assignment Model (CARAM). An illustration of the pre-processing stage for each MSA appears in Figure 9.

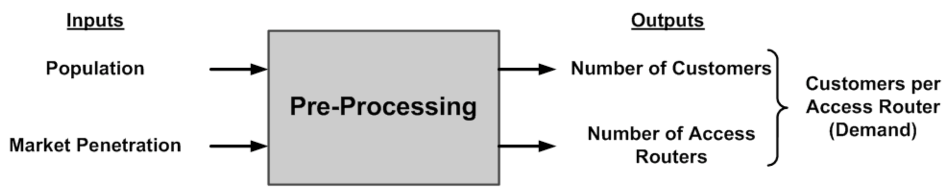

Figure 9. Pre-Processing Stage. This stage is applied to each MSA in turn. Inputs are the population and market penetration of each MSA. The outputs are the number of customers and access routers at each node.

1. Customer and Access Router Assignment Model (CARAM)

We consider a two-step deterministic model. The first step calculates the *number of customers* at a node based on the node's population and market penetration. We model the number of customers, c_i, at node i as

$$c_i = \lceil p_i w_i \rceil \tag{1.1}$$

where p_i is the population at node i, w_i is the (exogenously given) market penetration at node i, and $\lceil \bullet \rceil$ represents the ceiling operator.

The second step calculates the *number of access routers* at node i, a_i, based on the number of customers at node i, as:

$$a_i = \begin{cases} 1 & \text{if } c_i \le f_s \\ \left\lceil \dfrac{c_i}{f_m} \right\rceil & \text{if } c_i > f_s \end{cases} \tag{1.2}$$

where f_s is the maximum number of customer that a single access router can support and f_m is the maximum number of customers that multiple access routers can support. We assume that $f_m \le f_s$.

We illustrate the behavior of CARAM as a function of the number of customers in Figure 10. We assume the number of customers, c_i, in each node is linearly proportional to the customer population. Following equation (1.2), the number of access routers, a_i, is an increasing step function of c_i with steps occurring on a regular interval except for the first and second step. We also show *customers per access router*, denoted b_i, to illustrate the effect of parameters f_s and f_m. If the c_i is less then f_s, then b_i is bounded above by f_s. Otherwise, it is bounded above by f_m. The number of customers per access router is a discrete step function. As an example, given $f_s = 60$ and $f_m = 40$, a node with 500 customers would have an assumed 13 access routers with 38 customers per access router.

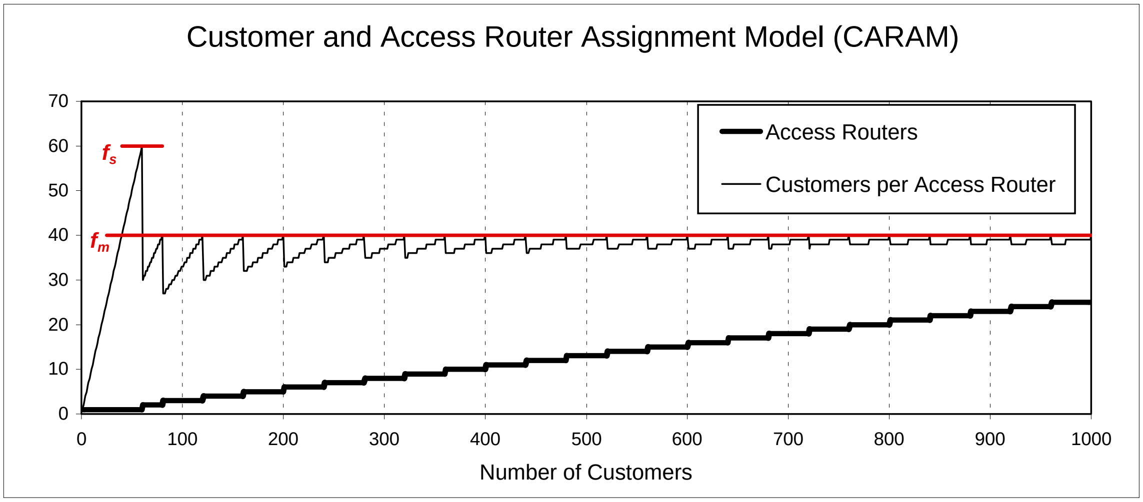

Figure 10. The Customer and Access Router Assignment Model (CARAM) prescribes the number of customers and access routers at a node given the population and market penetration at the node. The number of customers at a node is linearly proportional to the weighted population. The number of access routers assigned is dependent upon the number of customers. (i.e., a node with 500 customers would have 13 access routers with 38 customers per access router)

B. BACKBONE TOPOLOGY GENERATION

In the Backbone Topology Generation Stage, we interconnect the nodes associated with each geographic market into one network. We use the number of customers and number of access routers for each node (from the pre-processing stage) along with the node locations as inputs to this stage. The *number of backbone routers* for each node and a set of *backbone links* (node-node links) are outputs. Together these form the *backbone topology*. We illustrate the inputs and outputs in Figure 11.

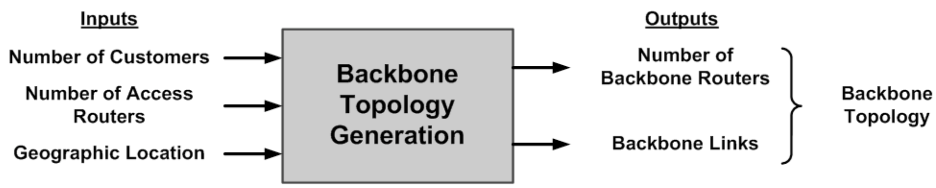

Figure 11. Backbone Topology Generation Inputs and Outputs.

We develop three topology generation models for this stage, one heuristic and two based on optimization models. We refer to the heuristic model as the Backbone Router and Link Assignment Model (BRLAM). We refer to the optimization models as the Minimum Cost Model (MCM) and the Maximum Flow Model (MFM). Both are mixed integer linear programs (MIPs).

1. Heuristic Backbone Router and Link Assignment Model

To generate a heuristic backbone topology, we first calculate the number of backbone routers at each node and then determine a set of links to connect the nodes.

a. Backbone Router Assignment Model (BRAM)

As discussed in Chapter II, backbone routers appear in pairs for redundancy reasons. We therefore model the number of backbone routers, b_i, at node i as

$$b_i = \begin{cases} 0 & \text{if } 0 \le a_i \le g_1 \\ 2 & \text{if } g_1 < a_i \le g_2 \\ 4 & \text{if } g_2 < a_i \le g_3 \\ 6 & \text{if } g_3 < a_i \end{cases} \tag{1.3}$$

where g_1, g_2, and g_3 are constant parameters satisfying $0 < g_1 \le g_2 \le g_3$. The behavior of the Backbone Router Assignment Model is illustrated in Figure 12.

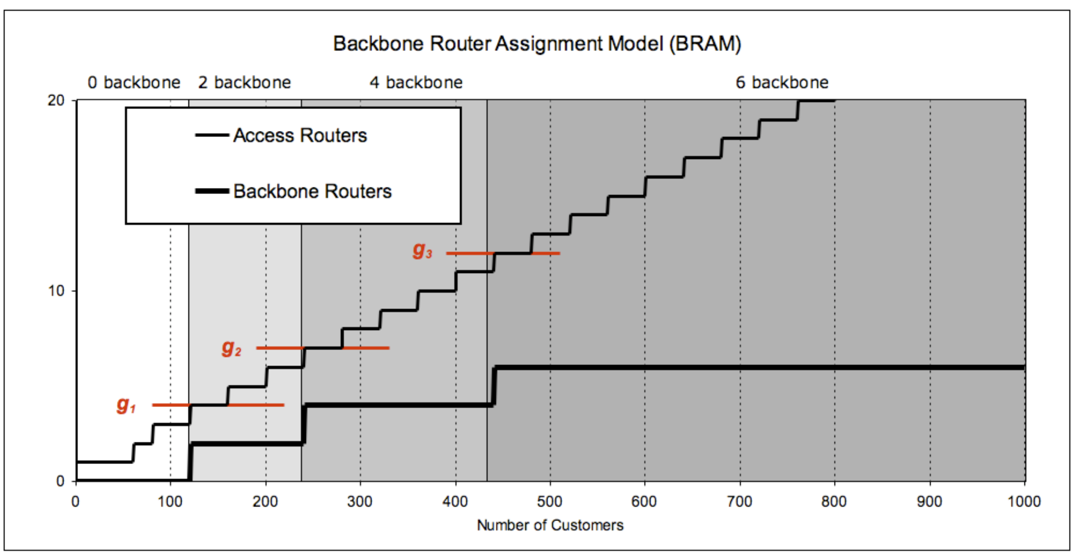

Figure 12. Backbone Router Assignment Model. $g_1 = 4$, $g_2 = 7$, $g_3 = 12$.

b. *Backbone Link Assignment Model (BLAM)*

Our heuristic topology generation model selects backbone links to connect the backbone nodes into a network. As nodes are connected, they become part of the backbone topology. We represent the backbone topology by the graph $G(N,A)$ where N is the set of nodes and A is the set of directed arcs in the backbone topology. We use a pair of directed arcs to represent each bidirectional link. We add arcs to A in four successive stages. The first stage involves connecting nodes with large b_i to each other, and in successive stages nodes with smaller b_i values are connected to the existing and growing network. We begin by partitioning the set of all nodes N into two subsets: C, the set of all core nodes (nodes with backbone routers) and E, the set of all edge nodes (nodes without backbone routers). We further partition C into three additional subsets: C_1, C_2, and C_3. We now have a partition of N into four subsets:

$$C_1 \cup C_2 \cup C_3 \cup E = N \tag{1.4}$$

We define the parameter $\lambda \in \{2,4,6\}$ to control the partition of the core nodes such that,

$$\begin{aligned}
C_1 &\equiv \{i \in N \mid b_i \geq \lambda\} \\
C_2 &\equiv \{i \in N \mid b_i = 4 \text{ and } i \notin C_1\} \\
C_3 &\equiv \{i \in N \mid b_i = 2,\ i \notin C_1, \text{ and } i \notin C_2\} \\
E &\equiv \{i \in N \mid b_i = 0\}
\end{aligned}$$

The arcs connecting nodes in C_j will be added in the jth iteration $(j = 1,2,3)$ and arcs connecting nodes in E will be added in the 4th iteration. We begin with $A = \{\varnothing\}$.

In the first stage, we use a procedure of link elimination based on triangles to choose links between nodes in C_1. For each combination of three nodes in C_1, we connect them to form a triangle, and let d_1, d_2, and d_3 represent the lengths of the legs in descending order ($d_1 \geq d_2 \geq d_3$). Then for some fixed choice of $\alpha \in [1.0, 2.0]$, if

$$\alpha d_1 \geq d_2 + d_3, \qquad (1.5)$$

we eliminate the link associated with the longest leg. We illustrate this procedure in Figure 13. If $\alpha = 1.0$, the longest legs will never be eliminated, and if $\alpha = 2.0$, the longest legs will always be eliminated. Finally, we add the arcs that represent the remaining links to A.

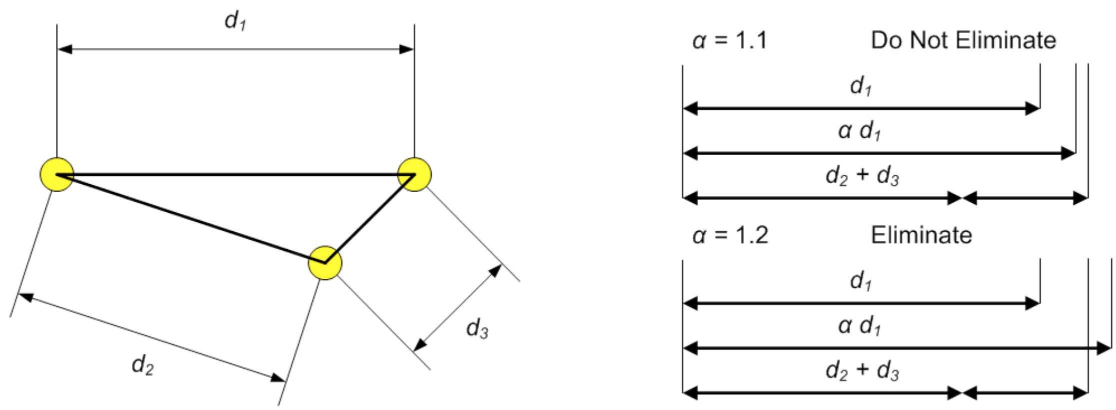

Figure 13. Link Elimination Procedure. We consider each combination of three nodes. The links between the nodes form the legs of a triangle. If the length of the longest leg of the triangle, multiplied by the parameter $\alpha \in [1.0, 2.0]$, is greater than the sum of the lengths of the two shortest legs, then eliminate the link associated with the longest leg.

In the second stage, we connect the nodes in C_2. We choose links such that each node in C_2 connects to the two nearest nodes in C_1 and we add the appropriate arcs to A.

In the third stage, we connect the nodes in C_3. We choose links such that each node in C_3 is connected to the two nearest nodes among C_1, C_2, and C_3 and we add the appropriate arcs to A. Note that in this stage, nodes in C_3 can be connected to other nodes in C_3.

In the fourth stage, we connect the nodes in E. We choose links such that each node in E is connected to the nearest node among C_1, C_2, and C_3 and we add the appropriate arcs to A.

The parameters λ and α have significant impact on the design of the backbone topology. We illustrate this impact in Figure 14. and in Figure 15.

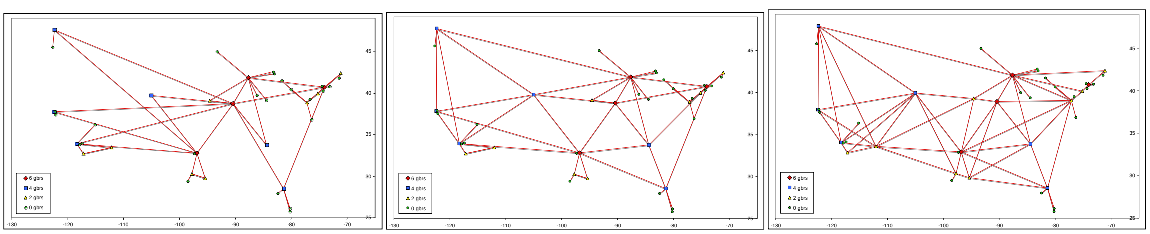

Figure 14. Networks generated with different values of λ.
(a) $\alpha = 1.06$ $\lambda = 6$, (b) $\alpha = 1.06$ $\lambda = 4$ (c) $\alpha = 1.06$ $\lambda = 2$

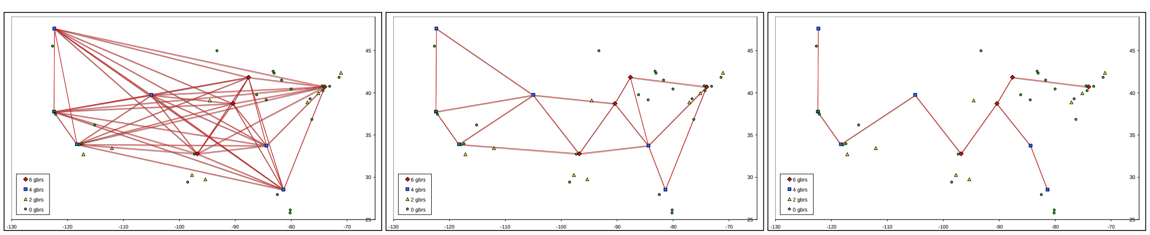

Figure 15. First layer core networks generated with different values of α.
(a) $\alpha = 1.00$ $\lambda = 4$, (b) $\alpha = 1.20$ $\lambda = 4$ (c) $\alpha = 2.00$ $\lambda = 4$

Due to the impact of parameters α and λ selecting their values is an important consideration.

2. Optimal Backbone Topology Models

The backbone topology design problem (BTDP) has three competing objectives: (1) to minimize cost; (2) to maximize flow; and (3) to be robust in terms of throughput capacity in the presence of link and/or node failure. The capacity and robustness objectives are counter to the cost objective. To increase either one, additional network components must be added, resulting in an increased cost. We will use goal based mixed integer programming to address this design problem. We formulate two mixed integer linear programs (MIPs), one that maximizes flow subject to a budget goal and a second that minimizes cost subject to a minimum flow goal. We implement robustness within each model via feasibility constraints.

The BTDP answers two questions. First, how many backbone routers should we place at each node? Backbone routers occur in pairs based on our design motif and thus our choice is among zero, one, two, or three pairs. Therefore, we have four types nodes corresponding to the number of backbone router pairs present. The backbone router configuration within each node type is deterministic. Thus, the cost of each node type is a function of the individual router and link costs. Likewise, the node type capacities are a function of the individual router capacities. Because the backbone routers and the inter-node links occur in pairs and the structure is symmetric, we calculate the node capacities using only one of the routers in each pair. We illustrate these relationships in Figure 16.

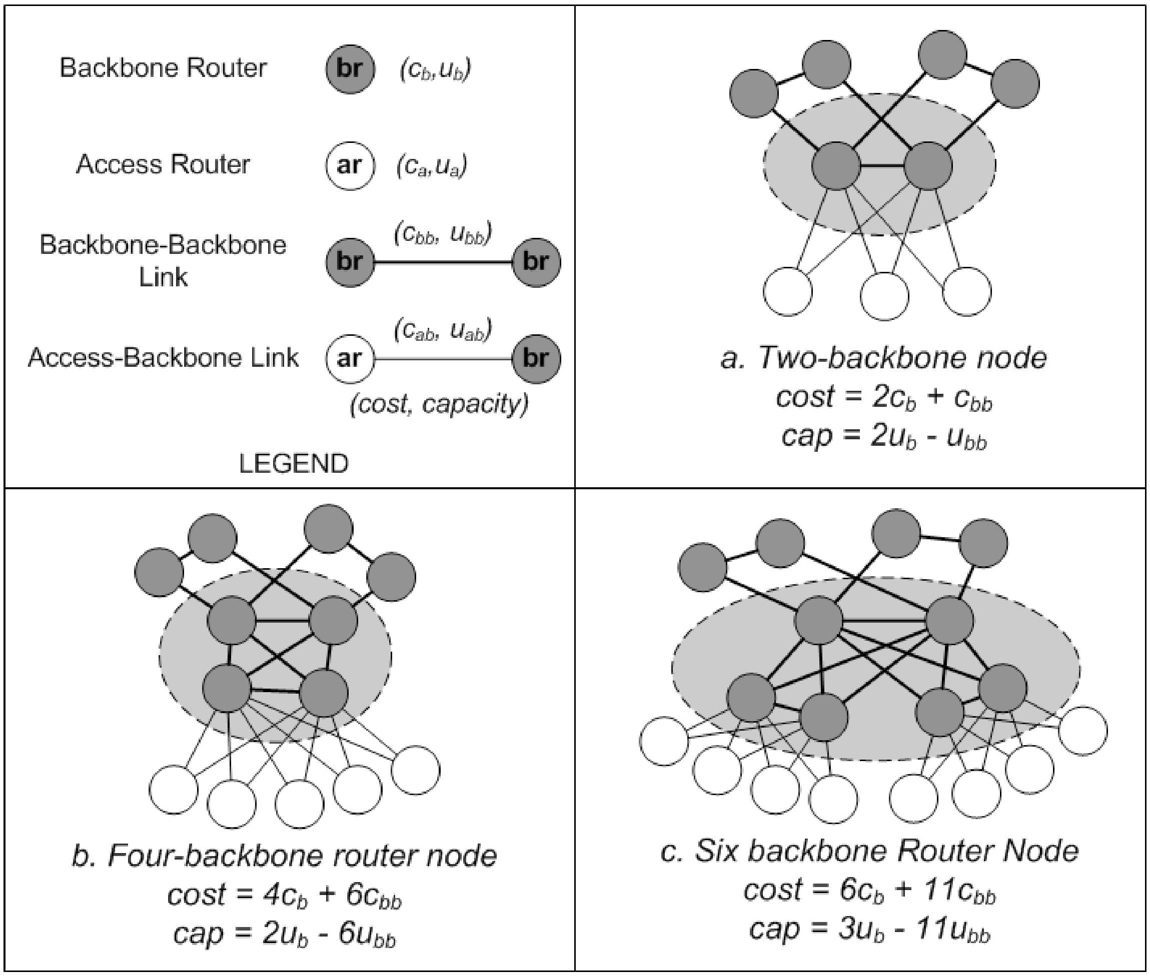

Figure 16. Cost and Capacity Assumption for each Node Type. The internal structure for each node follows directly from the design motifs observed in AS 7018 and illustrated in Figure 6.

The second question addressed by the BTDP is, which backbone topology links should be used to connect the nodes together? A potential backbone topology link exists between every pair of nodes in the network. We represent each of these bi-directional links by a pair of directed arcs. We classify each node as a core node or edge node, depending on whether or not it has backbone routers. Therefore, we have three types of arcs depending upon the core/edge classification of each arc's tail and head nodes. Edge-Edge links are precluded by construction. Each arc type has an associated cost and capacity, which is a function of its head and tail nodes. We illustrate these relationships in Figure 16. We allow for null backbone topology arcs as a fourth arc type; they have no capacity or cost. We list the backbone topology arc types in Table 4.

Table 4. Optimal Backbone Topology Model Arc Types

	Node Classification	
Arc Type	**Tail**	**Head**
0	na	na
1	Edge	Core
2	Core	Edge
3	Core	Core

We define the following indices, sets, parameters, and decision variables, to describe the backbone topology.

Index Use and Sets

i node; alias(j); $i \in N$

(i, j) arc; $(i, j) \in A$

p arc type; $p \in P = \{0, 1, 2, 3\}$

g node type; $g \in G = \{0, 2, 4, 6\}$

Parameters

a_i number of access routers at node i

u_g capacity of node of type g

v_p capacity of arc type p

Decision Variables

G_i^g binary variable equal to 1 if node i has g backbone routers, 0 otherwise.

H_i binary variable equal to 1 if node i is a core node, 0 otherwise.

E_{ij}^p binary variable equal to 1 if arc (i, j) is of type p, 0 otherwise.

A feasible region for the backbone topology, which is consistent with a hub and spoke design motif, is defined by the following system of equations, Υ.

Formulation of Backbone Topology Feasible Region

$$
\Upsilon = \left\{
\begin{array}{lll}
\sum_{g \in G} G_i^g = 1 & \forall i \in N & \text{(A1)} \\
G_i^0 = 1 - H_i & \forall i \in N & \text{(A2)} \\
E_{ij}^0 \le 1 & \forall (i,j) \in A & \text{(A3)} \\
E_{ij}^1 \le H_i & \forall (i,j) \in A & \text{(A4)} \\
E_{ij}^1 \le 2 - H_i - H_j & \forall (i,j) \in A & \text{(A5)} \\
E_{ij}^2 \le H_j & \forall (i,j) \in A & \text{(A6)} \\
E_{ij}^2 \le 2 - H_i + H_j & \forall (i,j) \in A & \text{(A7)} \\
E_{ij}^3 \le H_i & \forall (i,j) \in A & \text{(A8)} \\
E_{ij}^3 \le H_j & \forall (i,j) \in A & \text{(A9)} \\
\sum_{p} E_{ij}^p = 1 & \forall (i,j) \in A & \text{(A10)} \\
E_{ij}^0 = E_{ji}^0 & \forall (i,j) \in A & \text{(A11)} \\
E_{ij}^1 = E_{ji}^2 & \forall (i,j) \in A & \text{(A12)} \\
E_{ij}^3 = E_{ji}^3 & \forall (i,j) \in A & \text{(A13)} \\
\sum_{j|(i,j) \in A} \left(a_j v_1 E_{ij}^1 + v_3 E_{ij}^3 \right) + a_i v_1 H_i \le \sum_{g \in G} u_g G_i^g & \forall i \in N & \text{(A14)} \\
\sum_{j|(i,j) \in A} E_{ij}^1 \le 1 & \forall i \in N & \text{(A15)} \\
\sum_{j|(i,j) \in A} E_{ij}^3 \ge 2H_i & \forall i \in N & \text{(A16)} \\
G_i^g \in \{0,1\} \quad \forall i \in N, \forall g \in G & & \\
H_i \in \{0,1\} \quad \forall i \in N & & \\
E_{ij}^p \in \{0,1\} \quad \forall (i,j) \in A, \forall p \in P & &
\end{array}
\right.
$$

Constraint (A1) requires that every node can be of only one type. Constraint (A2) requires that any node with backbone routers is a core node. Constraint (A3) makes it is feasible for every arc to be a null arc (type 0). Constraints (A4) and (A5) require that a core-edge arc (type 1) is feasible between any two nodes if and only if the tail is a core node and the head is an edge node. Constraints (A6) and (A7) require that an edge-core arc (type 2) is feasible between any two nodes if and only if the tail is an edge node and the head is a core node. Constraints (A8) and (A9) require that a core-core arc is feasible only between a pair of core nodes. Constraint (A10) requires that every arc must be assigned a type and can only be of one type. Equations (A11), (A12), and (A13) require arc symmetry. Constraint (A14) requires the node capacity. A node can support as many outgoing arcs such that the sum of the outgoing arc capacities is less the node's capacity. The core-edge arc capacities are a multiple of the number of access routers in the edge node. Constraint (A15) requires that an edge node will only connect to one other node. Constraint (A16) requires that core nodes must have connections to at least two other core nodes.

Given a feasible backbone topology, the BTDP reduces to a multi-commodity network flow problem, were each pair of nodes in the network forms a source to destination (s-t) pair. Nodes in the network communicate under a gravity flow model, where the traffic between each s-t pair is proportional to the product of the number of customers at each node and a constant of proportionality.

Consider the following additional indices, parameters, and variables.

Index Use

i	node; alias(s,t); $i \in N$

Sets

R	set of all return arcs

Parameters

b_s	number of customers at node s
c_g	cost of node type g
d_{ij}	distance from node i to node j
e_p	cost per unit distance of arc of type p
f_p	fixed cost of using arc of type p
$budget$	maximum allowed cost
$flow$	minimum flow goal

Decision Variables

ρ	traffic scale parameter
X_{ij}^t	flow on arc (i,j) with destination t
Z_{ts}	flow on return arc (t,s)

The formulation of the Maximum Flow Model is as follows.

Maximum Flow Model Formulation

$$\max \quad \sum_{(s,t)\in R} Z_{ts} \tag{B1}$$

$$\begin{aligned} s.t. \quad & \sum_{i\in N, g\in G} c_g G_i^g + 2\sum_i a_i f_1 H_i + \sum_{(i,j)\in A} a_j \left(f_1 + d_{ij} e_1\right) E_{ij}^1 \\ & + \sum_{(i,j)\in A} a_i \left(f_2 + d_{ij} e_2\right) E_{ij}^2 + \sum_{\substack{(i,j)\in A \\ p\in\{0,3\}}} \left(f_p + d_{ij} e_p\right) E_{ij}^p \le budget \end{aligned} \tag{B2}$$

$$\sum_{t\in N} X_{ij}^t \le v_0 E_{ij}^0 + a_j v_1 E_{ij}^1 + a_i v_1 E_{ij}^2 + v_2 E_{ij}^3 \qquad \forall (i,j)\in A \tag{B3}$$

$$\sum_{j|(i,j)\in A} X_{ij}^t - \sum_{j|(i,j)\in A} X_{ji}^t = \begin{cases} Z_{ti}, & \text{if } i \ne t \\ -\sum_s Z_{ts}, & \text{if } i = t \end{cases} \qquad \forall i\in N, \forall t\in N \tag{B4}$$

$$Z_{ts} - \rho b_s b_t = 0 \qquad \forall (s,t)\in R \tag{B5}$$

$$X_{ij}^t \ge 0 \qquad \forall (i,j)\in A, \forall t\in N$$

$$Z_{ts} \ge 0 \qquad \forall (s,t)\in R$$

$$\rho \quad \text{URS}$$

$$G_i^g, H_i, E_{ij}^p \in \Upsilon$$

The objective function (B1) is the sum of the flows on all return arcs. The objective function value increases with the proportionality constant ρ.

Constraint (B2) enforces the budget. The first term accounts for the cost of a node based on its type. The second term accounts for the cost of connecting access routers within a hub node to the hub. The third term accounts for the cost of connecting access routers in non-hub nodes to hub nodes. The fourth term accounts for the cost of connecting hub nodes to other hub nodes. The sum of all the costs must be less then the budget.

Constraints (B3) through (B5) represent the multi-commodity flow model constraints. Constraint (B3) enforces the link capacity, equation (B4) enforces balance of flow at each node, and constraint (B5) enforces that source-destination flows between pairs of nodes will be proportional to the number of customers at each node.

The Minimum Cost Model is formulated as follows.

Maximum Flow Model Formulation

$$\begin{aligned} \min \quad & \sum_{i\in N, g\in G} c_g G_i^g + 2\sum_{i\in N} a_i f_1 H_i + \sum_{(i,j)\in A} a_j \left(f_1 + d_{ij} e_1\right) E_{ij}^1 \\ & + \sum_{(i,j)\in A} a_i \left(f_2 + d_{ij} e_2\right) E_{ij}^2 + \sum_{\substack{(i,j)\in A \\ p\in\{0,3\}}} \left(f_p + d_{ij} e_p\right) E_{ij}^p \end{aligned} \tag{C1}$$

$$s.t. \quad \sum_{(s,t)\in R} Z^{ts} \geq flow \tag{C2}$$

$$\sum_{t\in N} X_{ij}^t \leq v_0 E_{ij}^0 + a_j v_1 E_{ij}^1 + a_i v_1 E_{ij}^2 + v_2 E_{ij}^3 \qquad \forall (i,j) \in A \tag{C3}$$

$$\sum_{j:(i,j)\in A} X_{ij}^t - \sum_{j:(i,j)\in A} X_{ji}^t = \begin{cases} Z_{ti}, & \text{if } i \neq t \\ -\sum_{s\in N} Z_{ts}, & \text{if } i = t \end{cases} \qquad \forall i \in N, \forall t \in N \tag{C4}$$

$$Z_{ts} - \rho b_s b_t = 0 \qquad \forall (s,t) \in R \tag{C5}$$

$$X_{ij}^t \geq 0 \qquad \forall (i,j) \in R, t \in N$$

$$Z_{ts} \geq 0 \qquad \forall (s,t) \in R$$

$$\rho \quad \text{URS}$$

$$G_i^g, H_i, E_{ij}^p \in \Upsilon$$

The objective (C1) represents the cost of the network. Each term is the same as the terms in equation (B2). Constraint (C2) enforces that the total flow across return arcs must be greater then the flow goal. Constraints (C3) through (C5) represent the multi-commodity flow model constraints and are the same as constraints (B3) through (B5).

C. POST-PROCESSING: BUILDING A ROUTER-LEVEL MAP

In the Post-Processing Stage, we generate a router-level topology from the backbone topology. The router-level topology is deterministic and based on a design motif of a redundant hierarchical tree as described in Figure 16. Inputs to this stage are the number of access routers and backbone routers at each node, along with, the backbone topology links, which connect the nodes. We illustrate this stage in Figure 17. and Figure 18.

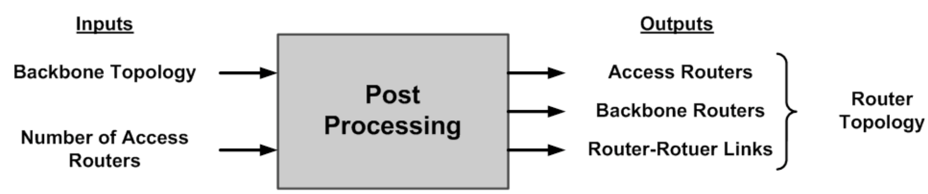

Figure 17. Post Processing Stage

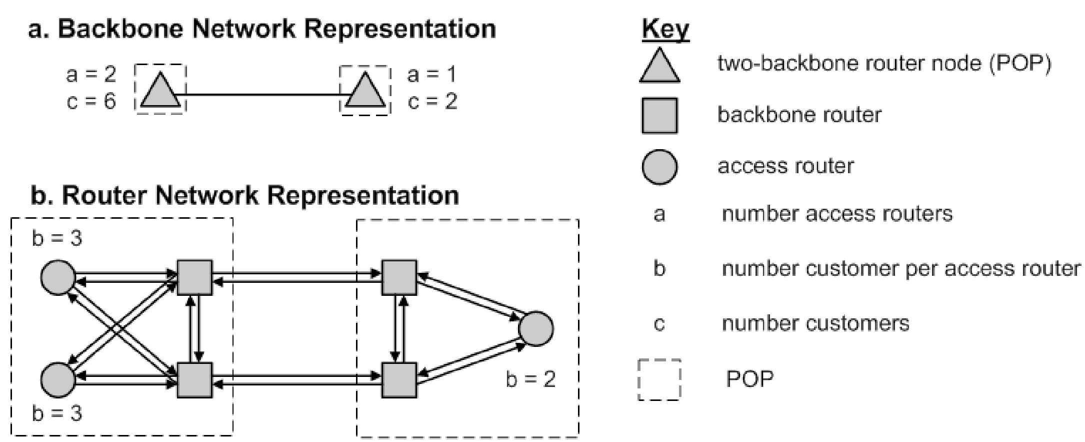

Figure 18. We build a router-level topology from the backbone topology in the post processing stage. a. Backbone representation with (2) two-backbone router core nodes. b. Equivalent router-level topology.

IV. ANALYZING TOPOLOGIES

In the previous chapters, we have analyzed an existing ISP network and identified relationships between both its structure and the assumed underlying customer population that it supports. Using these relationships, we have developed the means to generate backbone and router-level topologies for any collection of geographically dispersed customer populations. We have formulated three models for generating the backbone topology of the network, one using a heuristic method and two using optimal methods.

We now generate topologies using each of the backbone topology generation models developed in Chapter III. To allow easy comparison of the topologies, we use the following methodology. We first generate a topology using the heuristic generation model. We then use the cost and throughput of this topology as the budget and minimum flow constraints in the optimization-based generation models. Furthermore, we use the topology generated by the heuristic as an initial feasible solution in the optimization models. We compare the topologies using both the backbone and router representations. We illustrate this methodology in Figure 19.

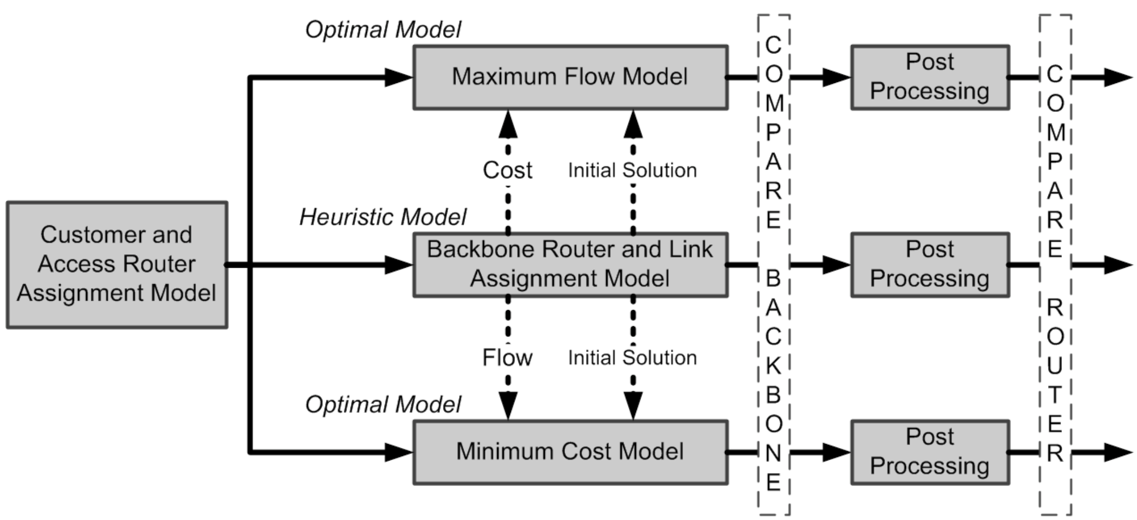

Figure 19. Analysis Methodology

A. ANALYSIS DATA SETS

We use the set of MSAs for AS 7018 as the input data for our topology generation and analysis. We select eight subsets of the MSA list to represent customer populations that range from regional (e.g., Southern California and Eastern United States) to national (e.g., the entire United States). In addition to the number of MSAs, we also try to capture different geometries, e.g., national network with many large MSAs (hub heavy) and national network with many small MSAs (spoke heavy).

A summary of the MSA subsets appears in Table 5. The full MSA data set and subsets are listed in the Appendix. We illustrate the MSA subsets in Figure 20. We represent the MSAs by dots that are proportional in size to the MSA's population.

Table 5. Metropolitan Statistical Area Subset Summary

Subset	**Number of MSAs**	**Description**
1	7	Small Network
2	10	Southern California
3	14	Chicago-Atlanta-New York
4	17	Western United States
5	52	Eastern United States
6	79	United States Edge Heavy
7	42	United States Core Heavy
8	89	All MSAs

We list the router and link cost and capacities used in the models in Table 6. We use fixed hardware costs based upon a recent Cisco pricing catalog (Cisco, 2003).

Throughout the remainder of this chapter, we use the terms MSA and node interchangeably. As before, *Core nodes* are nodes that have backbone routers and *edge nodes* are nodes that do not have backbone routers. *Equal Cost* refers to the Maximum Flow Model solution and *Equal Flow* refers to the Minimum Cost Model solution.

Figure 20. Metropolitan Statistical Area Subsets 1-8

Table 6. Model Cost and Capacity Parameters

Network Component	Capacity [Gps]	fixed [$K]	Per Mile [$K/mile]
Access Router	10	0	-
Backbone Router	150	125	-
Access-Backbone Link	1	15	1
Backbone-Backbone Link	10	350	5

1. Subset 1: Small Network

Subset 1 contains only 7 nodes. One core node has four backbone routers, while the others each have two. The core nodes are fully connected and the edge nodes each connect to one of the core nodes.

The three topologies appear in Figure 21. The equal cost topology is the same as the heuristic topology, while in the equal flow topology solution the four-backbone router node becomes a two-backbone router node. A constraint in the optimal models requires that each core node connect to at least two other core nodes. This constraint implies that a network must have at least three core nodes.

We list the numerical results of the Backbone Generation Models on subset 1 in Table 7.

Table 7. Subset 1 Results

	Cost (% Heuristic) [$K]		Flow (%Heuristic)[Gps]	
Heuristic	32,625		38.74	
Equal Cost	32,625	(100.0%)	38.74	(100.0%)
Equal Flow	30,685	(94.1%)	38.74	(100.0%)

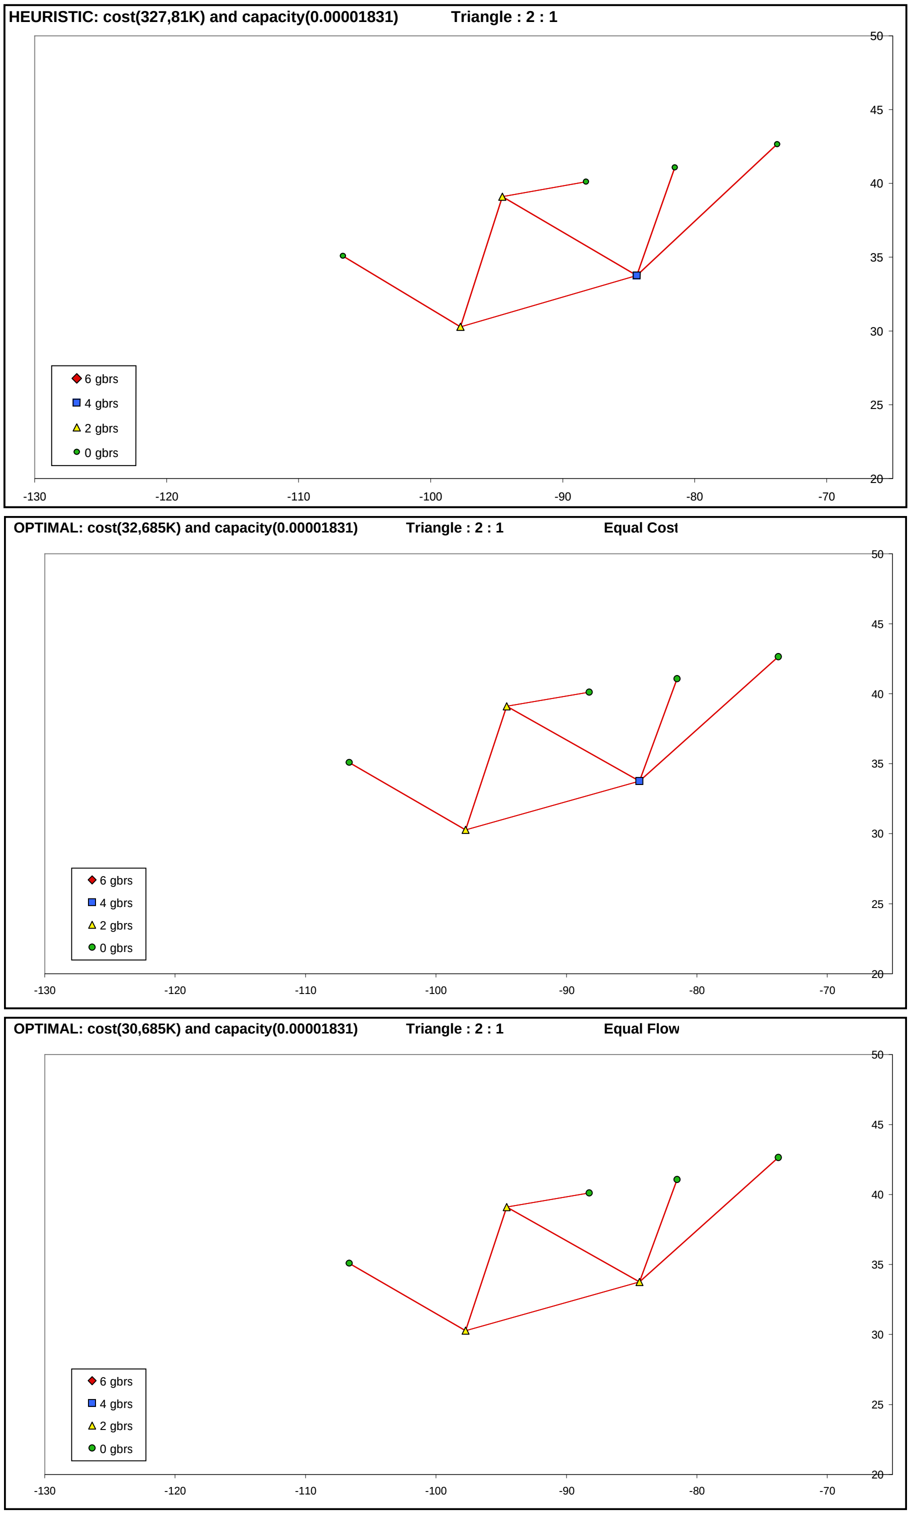

Figure 21. Subset 1 Backbone Topology Generation Solutions. a. Heuristic Model Solution. b. Optimal Maximum Flow Modes (Equal Cost) solution. c. Optimal Minimum Cost Model (Equal Flow) solution.

2. Subset 2: Southern California Region

Subset 2 represents a small regional area, specifically Southern California, Arizona and Nevada. The subset has 10 nodes. Most of the nodes are moderately sized and serve as core nodes in the heuristic solution. Three of the five core nodes have more than two backbone routers each.

The equal cost solution achieves considerably higher throughput by redistributing budget away from the large core nodes and then promoting all edge nodes to core nodes. This dramatically increases the capacity of all nodes and arcs throughout the network.

The equal flow topology solution downsizes the four- and six-backbone router core nodes to two-backbone router core nodes and eliminates one core-core link reducing the link structure to a loop.

We list the numerical results of the Backbone Generation Models on subset 2 in Table 8.

Table 8. Subset 2 Results

	Cost (% Heuristic) [$K]		**Flow (%Heuristic)[Gps]**	
Heuristic	30,384		32.24	
Equal Cost	30,413	(100.1%)	195.35	(605.8%)
Equal Flow	23,295	(76.7%)	32.24	(100.0%)

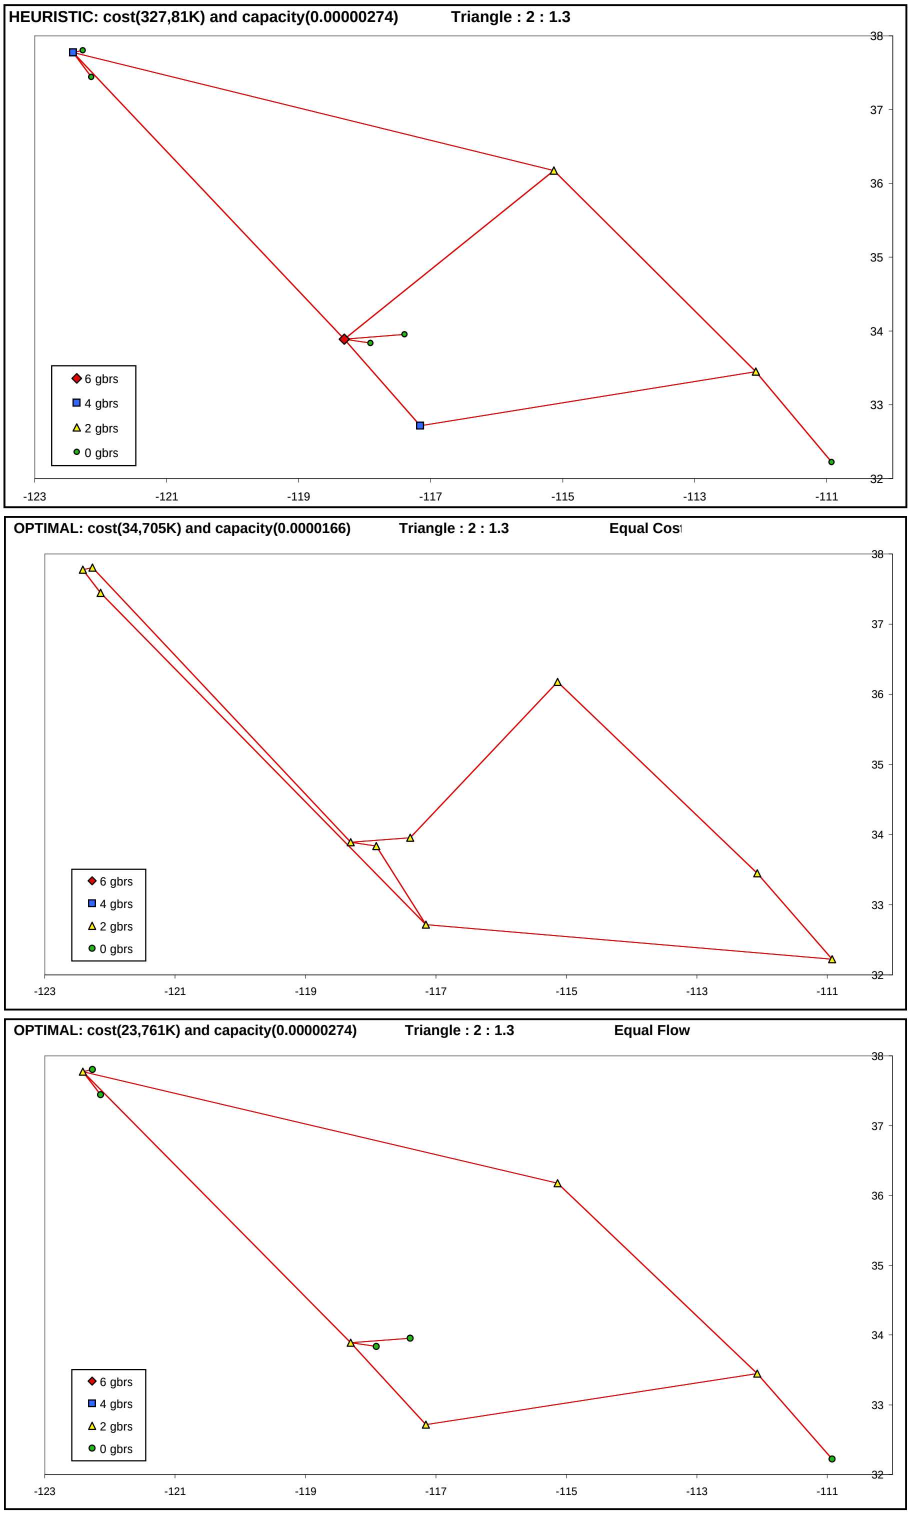

Figure 22. Subset 2 Backbone Topology Generation Solutions. a. Heuristic Model Solution. b. Optimal Maximum Flow Modes (Equal Cost) solution. c. Optimal Minimum Cost Model (Equal Flow) solution.

3. Subset 3: Three Large MSAs

Subset 3 represents a region with three large nodes surrounded by a handful of small nodes. The subset has 14 nodes total. The heuristic assigns backbone routers to each of the large nodes and no backbone routers to any of the small nodes. The core nodes are then fully connected into a triangle with the edge nodes connecting to the nearest core node.

In the equal cost solution, we find a similar redistribution of the infrastructure as in subset 2. Large core nodes are downsized and all but two edge nodes are promoted to core nodes. The core nodes are connected in a loop.

In the equal flow solution, we also find all of the large core nodes reduced and several of the edge nodes promoted. However, the core nodes are not connected in one loop but rather two small triangles linked by one long link.

We illustrate the three solutions in Figure 23. We list the numerical results of the Backbone Generation Models on subset 3 in Table 9.

Table 9. Subset 3 Results

	Cost (% Heuristic) [$K]		**Flow (%Heuristic)[Gps]**	
Heuristic	47,695		38.74	
Equal Cost	47,654	(99.9%)	158.04	(414.5%)
Equal Flow	37,715	(79.1%)	38.74	(100.0%)

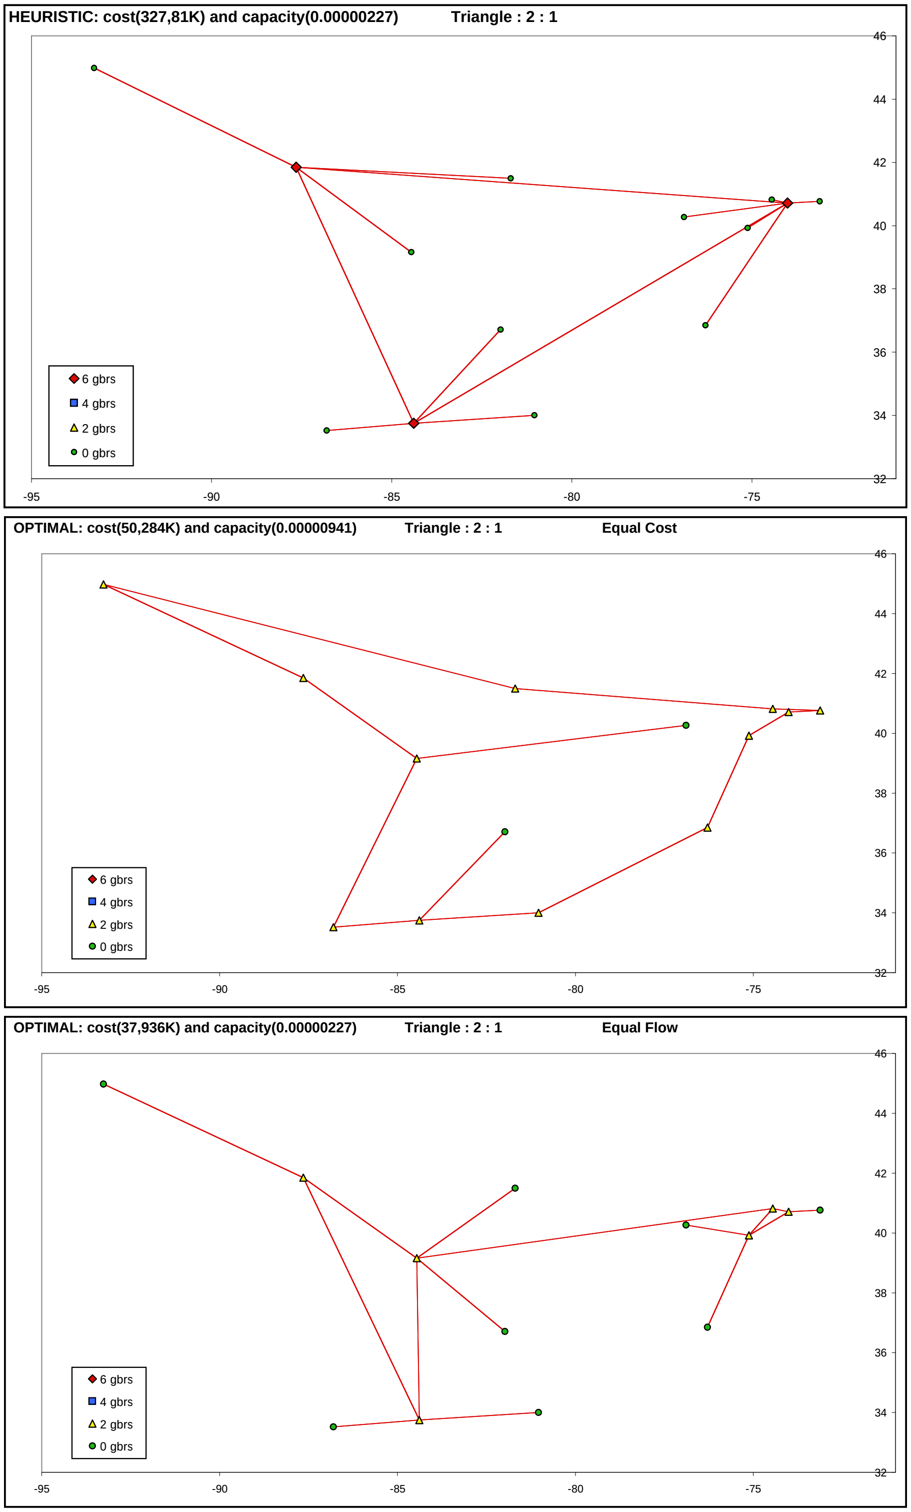

Figure 23. Subset 3 Backbone Topology Generation Solutions. a. Heuristic Model Solution. b. Optimal Maximum Flow Modes (Equal Cost) solution. c. Optimal Minimum Cost Model (Equal Flow) solution.

4. Subset 4: Western United States

Subset 4 represents the Western United States. It has 17 nodes. The heuristic assigns backbone routers to six of the nodes making them core nodes and then connects them in a loop with the edge nodes connecting to the nearest core node.

In the equal cost solution, we see the same pattern of the previous two subsets. In the equal flow solution, we find a simple reduction of all of the large core nodes to two-backbone router core nodes. No links are eliminated.

We illustrate the solutions in Figure 24. We list the numerical results of the Backbone Generation Models on subset 4 in Table 10.

Table 10. Subset 4 Results

	Cost (% Heuristic) [$K]		Flow (%Heuristic)[Gps]	
Heuristic	53,910		40.08	
Equal Cost	53,659	(99.5%)	126.22	(314.9%)
Equal Flow	45,910	(85.2%)	40.08	(100.0%)

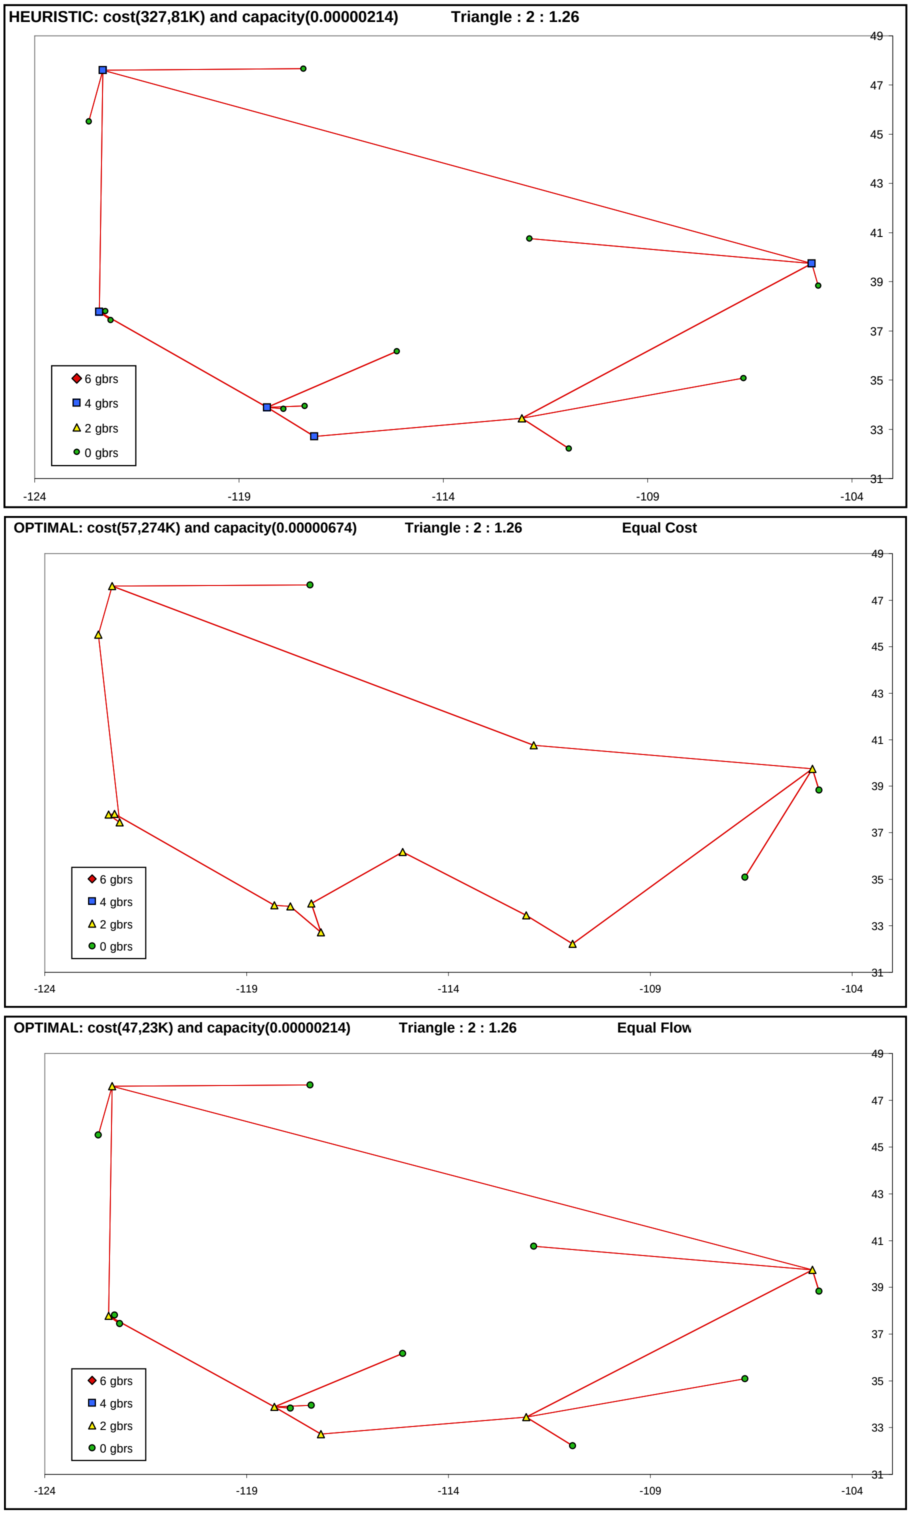

Figure 24. Subset 4 Backbone Topology Generation Solutions. a. Heuristic Model Solution. b. Optimal Maximum Flow Modes (Equal Cost) solution. c. Optimal Minimum Cost Model (Equal Flow) solution.

5. Subset 5: North Eastern United States

Subset 5 represents the North Eastern United States. It has 52 nodes the vast majority with small populations. The heuristic builds five core nodes. The three six and four backbone router core nodes are fully connected in a triangle and the two backbone router core nodes form a loop beginning an ending at one of the six backbone router core nodes. The edge nodes all connect to one of the core nodes.

Due to run time considerations, we implement an additional constraint in the equal cost and equal flow models for subsets 5, 6, 7, and 8. This constraint fixes the heuristic solution's edge nodes preventing them being upgraded to core nodes. For subset 5, we found no improvement in the equal cost solution's throughput.

In the equal cost solution, cost was improved by reducing all of the core nodes to two-backbone routers and changing core-core links to form a loop.

We illustrate the three solutions in Figure 25. We list the numerical results of the Backbone Generation Models on subset 5 in Table 11.

Table 11. Subset 5 Results

	Cost (% Heuristic) [$K]		Flow (%Heuristic)[Gps]	
Heuristic	92,899		75.0	
Equal Cost	92,899	(100.0%)	75.0	(100.0%)
Equal Flow	77,080	(83.0%)	75.0	(100.0%)

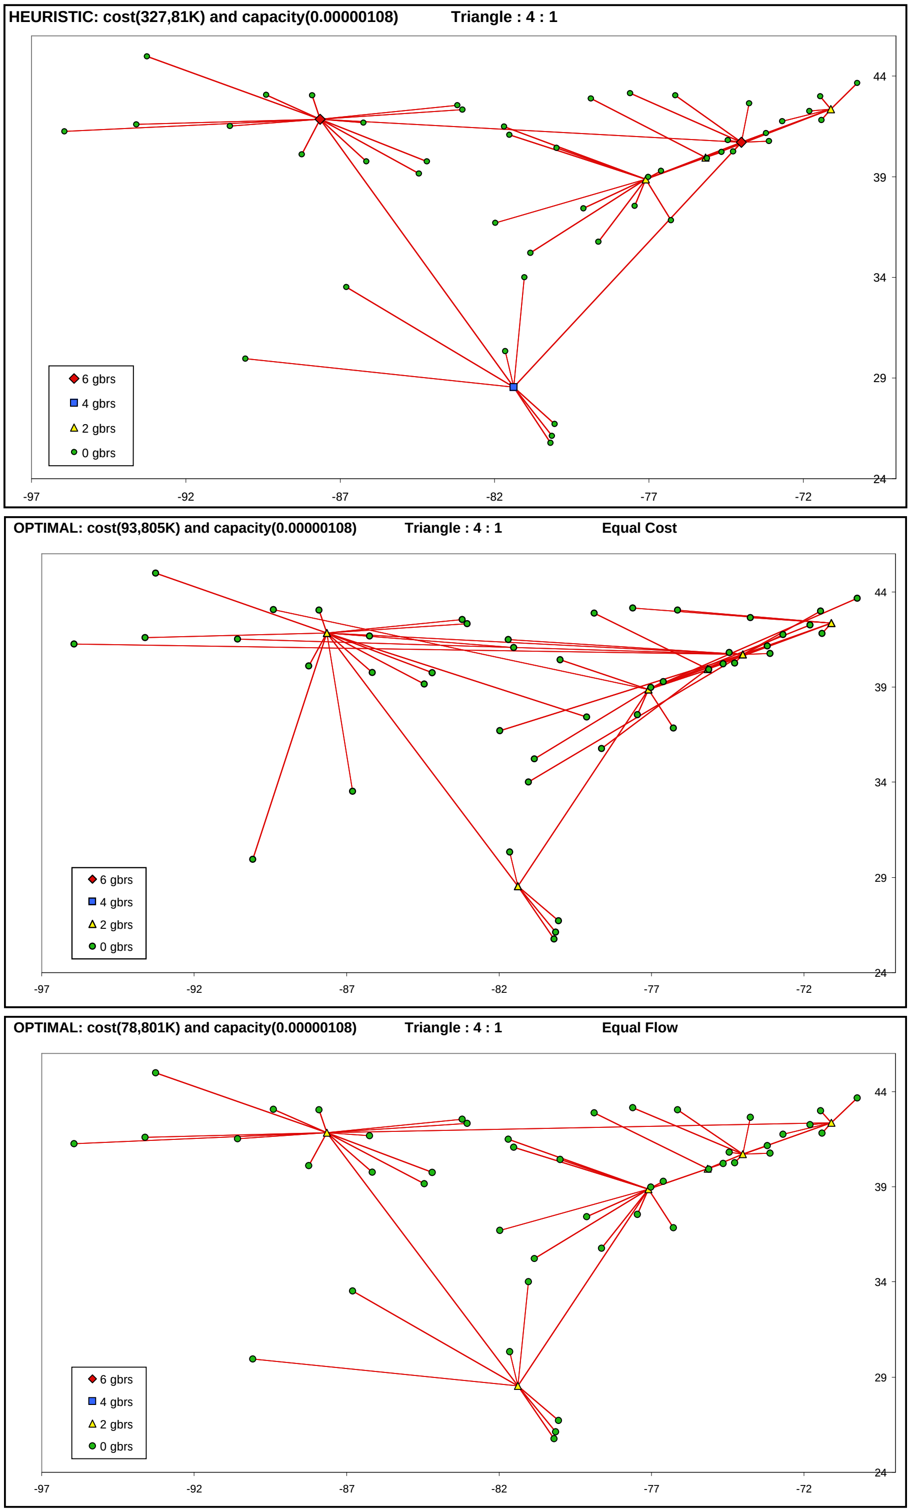

Figure 25. Subset 5 Backbone Topology Generation Solutions. a. Heuristic Model Solution. b. Optimal Maximum Flow Modes (Equal Cost) solution. c. Optimal Minimum Cost Model (Equal Flow) solution.

6. Subset 6: United States Edge Heavy

Subset 6 represents the United States with a large number of small MSAs. It has 79 nodes. The heuristic builds nine core nodes. The core nodes are connected by a mesh like pattern of links with the edge nodes connecting to the nearest core node.

The equal cost solution is identical to the heuristic solution due to the edge node restriction discussed in subset 5.

We still improve the cost with the equal flow solution by reducing all of the core nodes to two -backbone routers and changing core-core links to form a loop as in subset 5.

We illustrate the three solutions in Figure 26. We list the numerical results of the Backbone Generation Models on subset 6 in Table 12.

Table 12. Subset 6 Results

	Cost (% Heuristic) [$K]		Flow (%Heuristic)[Gps]	
Heuristic	256,355		130.99	
Equal Cost	256,355	(100.0%)	130.99	(100.0%)
Equal Flow	148,174	(57.8%)	130.99	(100.0%)

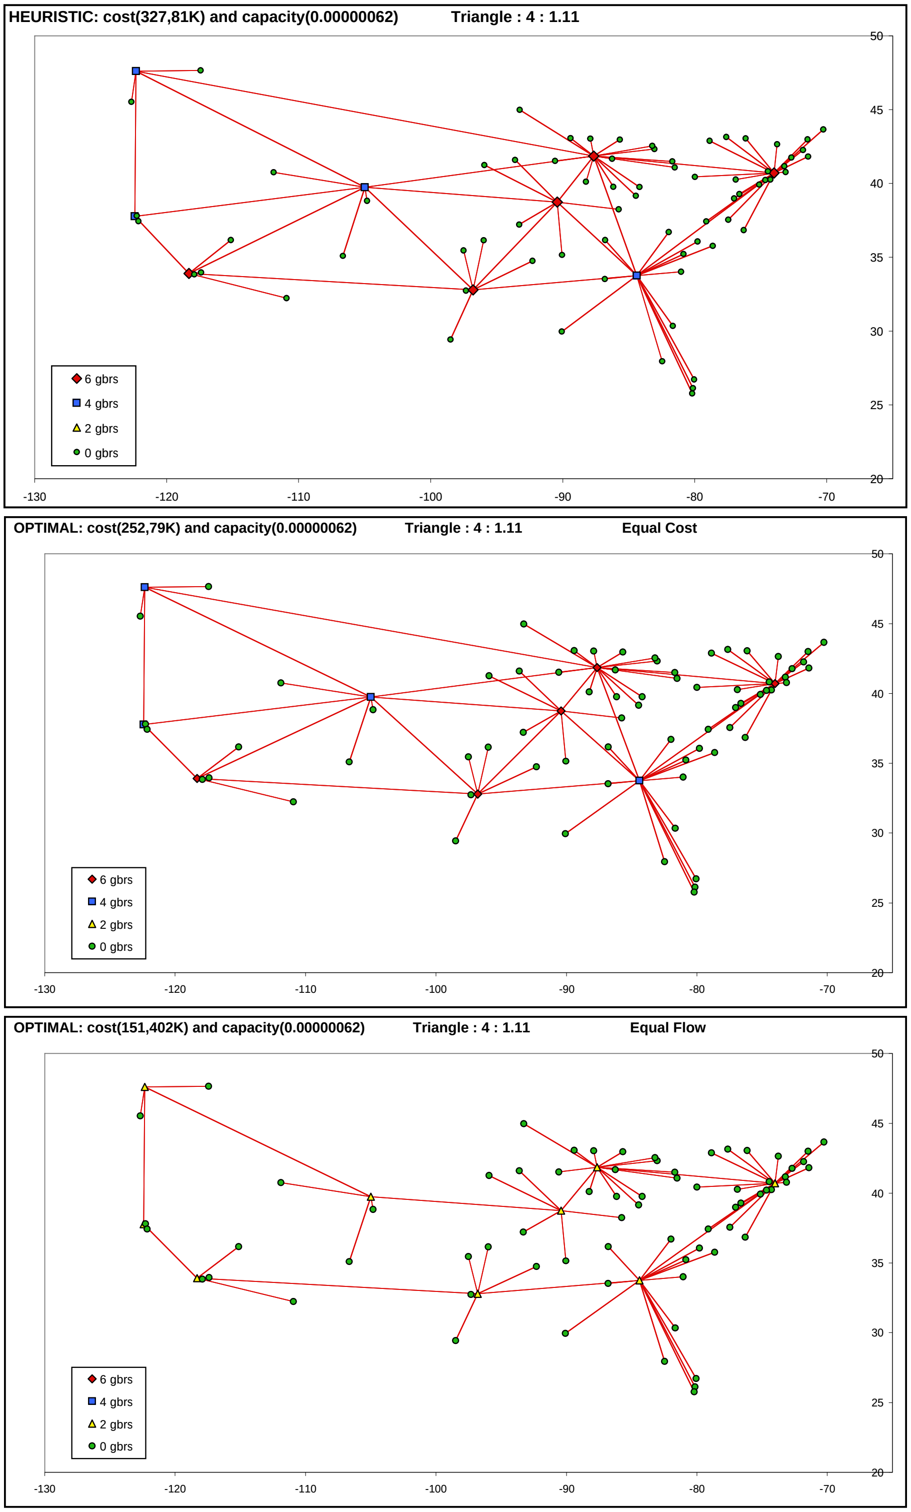

Figure 26. Subset 6 Backbone Topology Generation Solutions. a. Heuristic Model Solution. b. Optimal Maximum Flow Modes (Equal Cost) solution. c. Optimal Minimum Cost Model (Equal Flow) solution.

7. Subset 7: United States Core Heavy

Subset 7 represents the United States with only a few number of small MSAs. It has 42 nodes. The heuristic builds 18 core nodes. The core nodes are connected by a mesh like pattern of links with the edge nodes connecting to the nearest core node.

The equal cost solution is identical to the heuristic solution due to the edge node restriction discussed in subset 5.

We still improve the cost with the equal flow solution by reducing all of the core nodes to two-backbone routers and changing core-core links to form a loop as in subset 5.

We illustrate the solutions in Figure 27. We list the numerical results of the Backbone Generation Models on subset 7 in Table 13.

Table 13. Subset 7 Results

	Cost (% Heuristic) [$K]	Flow (%Heuristic)[Gps]
Heuristic	274,154	126.00
Equal Cost	274,154 (100.0%)	126.00 (100.0%)
Equal Flow	137,578 (50.2%)	126.00 (100.0%)

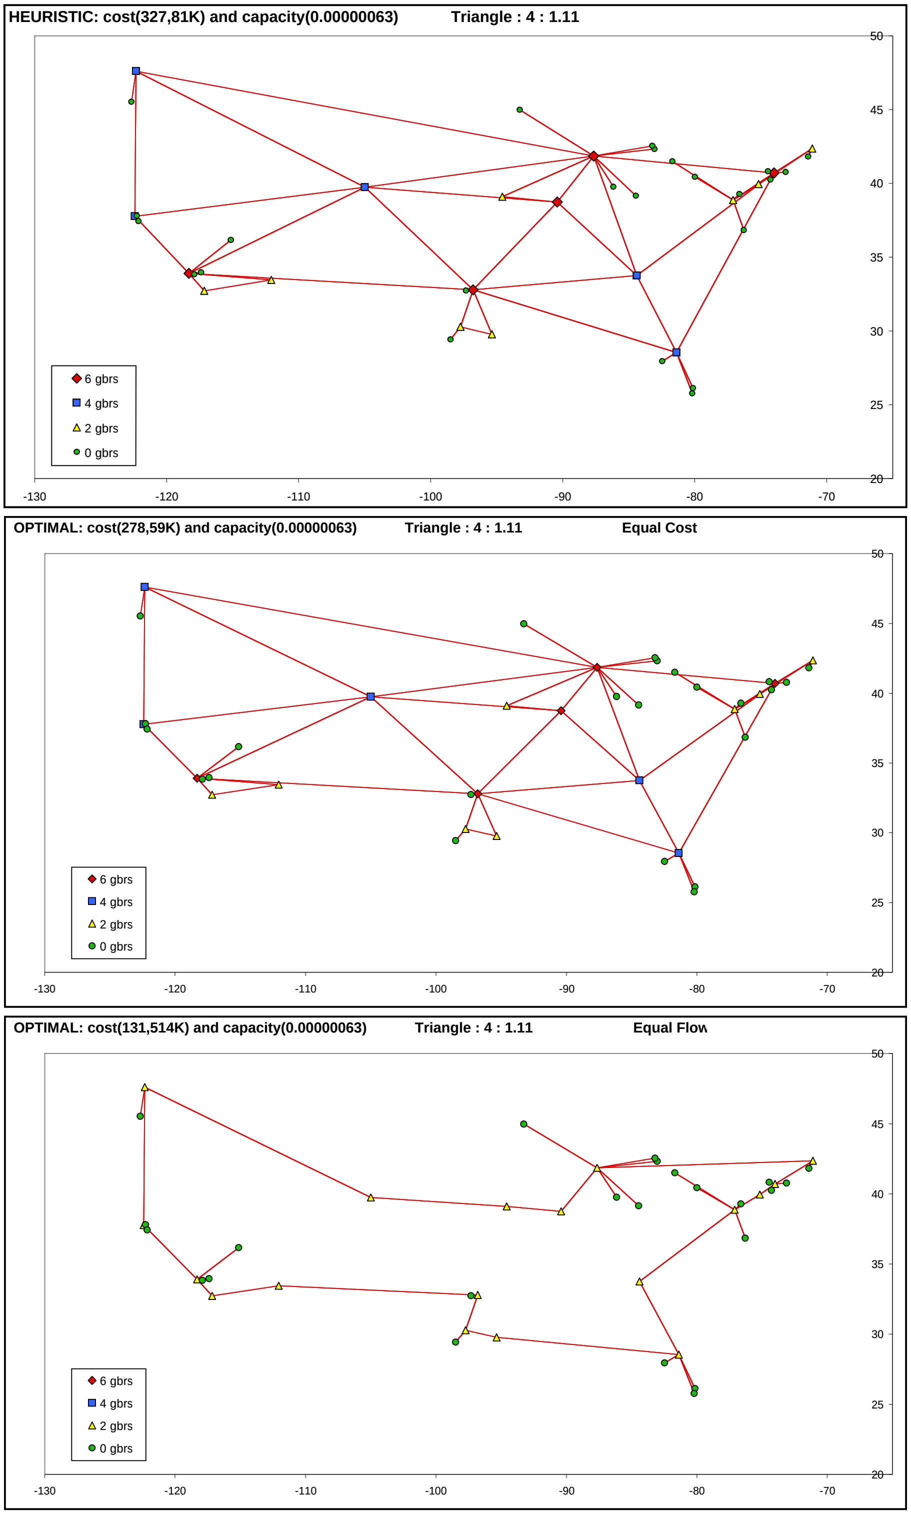

Figure 27. Subset 7 Backbone Topology Generation Solutions. a. Heuristic Model Solution. b. Optimal Maximum Flow Modes (Equal Cost) solution. c. Optimal Minimum Cost Model (Equal Flow) solution.

8. Subset 8: All MSAs

Subset 8 represents the United States including all of the MSAs. It has 89 nodes. The heuristic builds 18 core nodes. The core nodes are connected by a mesh like pattern of links with the edge nodes connecting to the nearest core node.

The equal cost solution is identical to the heuristic solution due to the edge node restriction discussed in subset 5.

We still improve the cost with the equal flow solution by reducing all of the core nodes to two-backbone routers and changing core-core links to form a loop as in subset 5.

We illustrate the solutions in Figure 28. We list the numerical results of the Backbone Generation Models on subset 8 in Table 14.

Table 14. Subset 8 Results

	Cost (% Heuristic) [$K]		Flow (%Heuristic)[Gps]	
Heuristic	302,221		169.61	
Equal Cost	302,221	(100.0%)	169.61	(100.0%)
Equal Flow	159,709	(52.8%)	169.61	(100.0%)

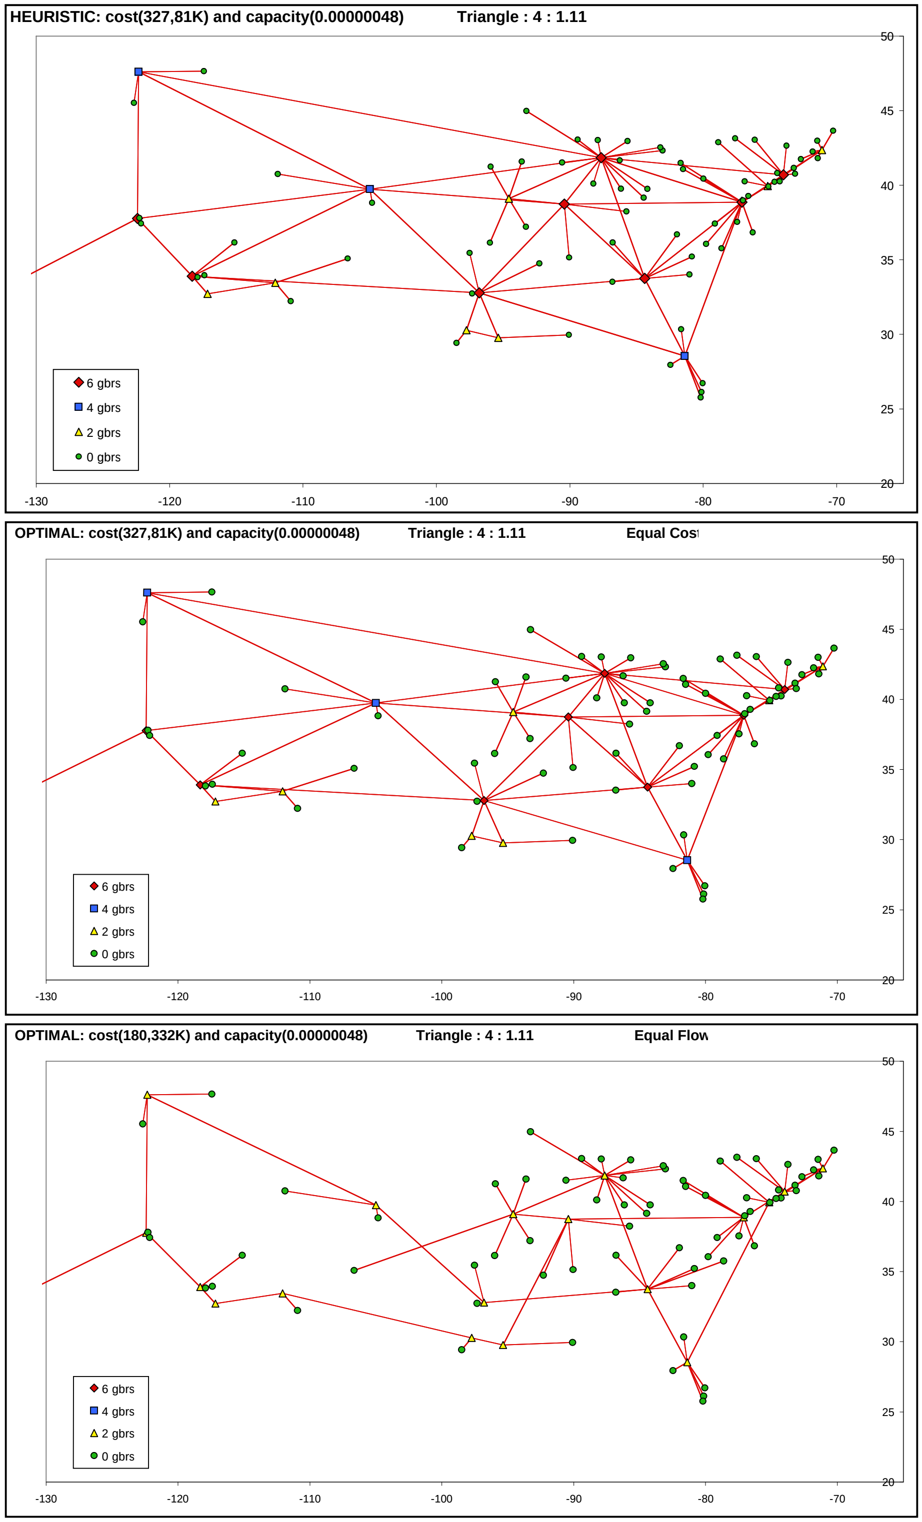

Figure 28. Subset 8 Backbone Topology Generation Solutions. a. Heuristic Model Solution. b. Optimal Maximum Flow Modes (Equal Cost) solution. c. Optimal Minimum Cost Model (Equal Flow) solution.

B. ANALYSIS RESULTS

The results of the backbone topology generation stage are interesting and informative. They indicate that backbone topology generation models behave as expected. In each case, the optimal models produced a solution at least as good as the heuristic model and for the most part improved upon it. However, the backbone topologies are abstractions of router-level topologies, which are of real interest to us. Therefore, for each case, we generate router-level topologies from the backbone topologies and using these, we reevaluate cost and total throughput. We illustrate the results in Figure 29.

The cost of the generated router-level topologies matches exactly the cost of the backbone topologies, and it increases with the size of the network. We list the number of router and arcs (two arcs per link) in Table 16.

The throughput of the router-level topologies follows a similar trend, increasing as the network size increases. However, the throughputs are not the same as the backbone representations. The backbone topology representation of a network ignores the router structure internal to nodes, and the backbone flow is based a maximum flow network model with no restrictions with regard to traffic engineering. We would expect then, the maximum flow on a backbone topology to be an upper bound on the maximum flow that the router-level topology could achieve. In our examples, this is not always the case. Many of the router network representations achieve higher throughputs then the backbone representations as seen in Table 15.

Table 15. Throughput achieved by the router topology representation relative to the backbone topology representation.

	Subset							
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Heuristic	100%	199%	356%	238%	100%	169%	254%	224%
Equal Cost	100%	37%	83%	73%	155%	169%	254%	224%
Equal Flow	100%	273%	211%	238%	143%	96%	98%	96%

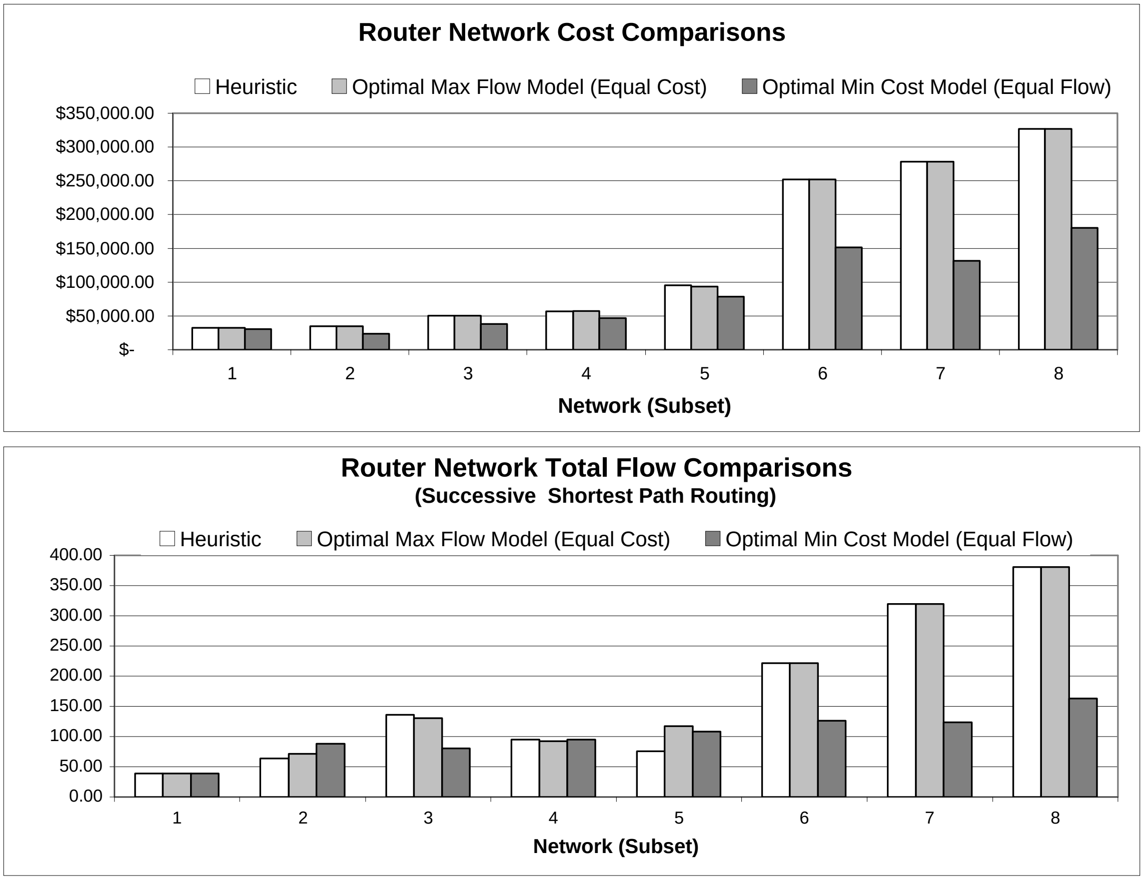

Figure 29. Router-level topology cost and flow comparisons for subsets 1-8. a. Cost comparison. b. Total throughput comparison using successive shortest path routing.

Table 16. Router and Arc Counts for Router-Level Topologies

	Subset							
Model	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Heuristic								
Total Routers	36	82	95	104	183	320	333	429
Total Arcs	140	330	386	414	738	1322	1386	1786
Equal Cost								
Total Routers	36	84	101	110	173	320	333	429
Total Arcs	140	320	380	412	684	1322	1386	1786
Equal Flow								
Total Routers	34	74	89	94	173	292	303	391
Total Arcs	130	286	348	364	680	1150	1176	1532

From a customer viewpoint, the total throughput capacity of the network is not of great concern. Rather, the ability of the network to deliver an expected level of bandwidth is more important. Therefore, for both shortest path and successive shortest path routing, we calculate the downstream customer bandwidth delivered by each router network when operating at maximum capacity. We assume that each customer expects 10 megabits per second of bandwidth (0.01 Gps). We illustrate the results in Figure 20 and Figure 31. Under single shortest path routing (naïve traffic engineering), the customers of the larger networks, do not receive the expected bandwidth. However, under successive shortest path routing (best case traffic engineering), the customers in every network receive the expected bandwidth. This illustrates the importance of traffic engineering and provides a secondary type of validation. The assumed parameters of our model (relative capacities) are reasonable and consistent with our design objectives.

We also consider router utilization, which is the fractional amount of a router's total throughput capacity that is used. For the eight subsets and three backbone generation models (under maxflow conditions), we illustrate access router utilization in Figure 32. and backbone router utilization in Figure 33. In all cases, backbone and access have considerable excess capacity indicating that the bottlenecks in the networks are links not routers.

We have evaluated the network topologies using two of three performance objectives, cost and throughput. We find that our heuristic produces topologies for which both the cost or the throughput can be improved upon using optimal methods. The third performance objective, robustness, we do not evaluate in this Thesis. Previous work by Barkley (2008) lays out a model for optimally attacking router-level topologies that follows in the spirit of Brown et al. (2006).

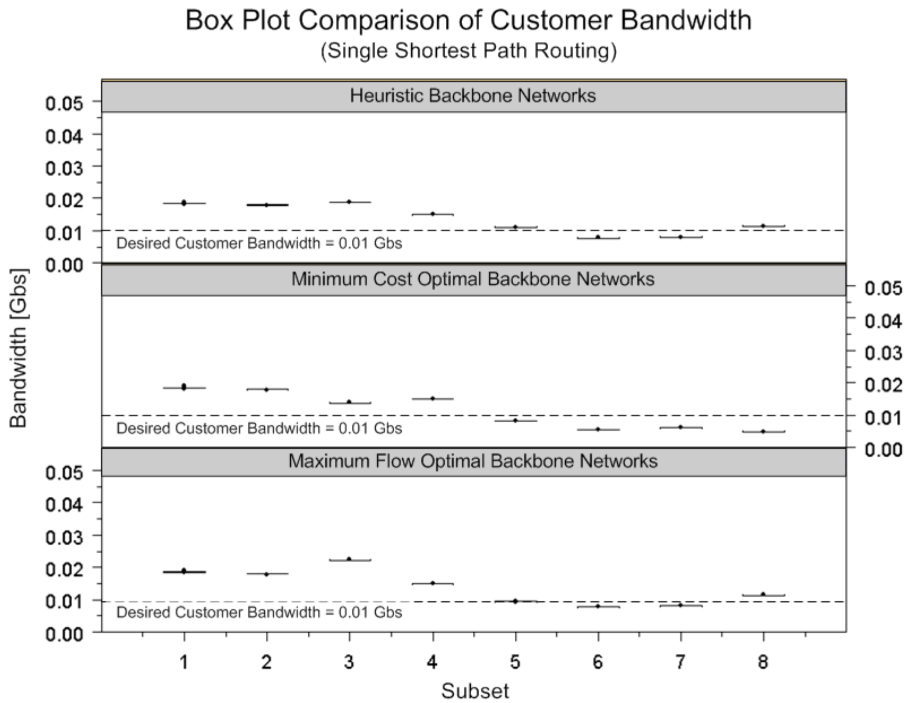

Figure 30. Achieved Customer Bandwidth (Shortest Path Routing)

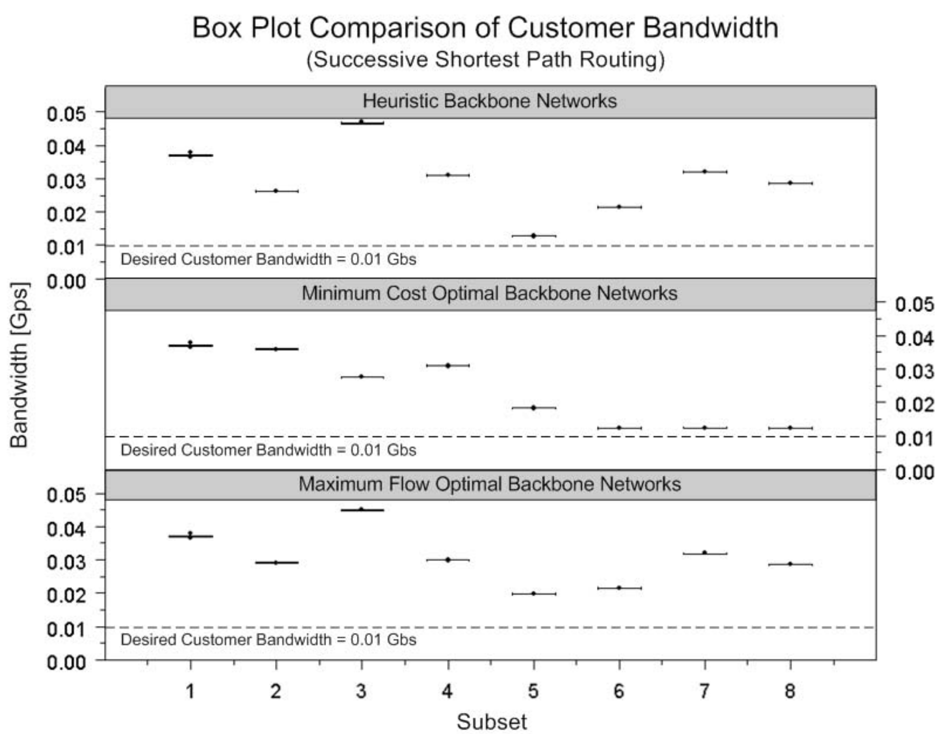

Figure 31. Achieved Customer Bandwidth (Successive Shortest Path Routing)

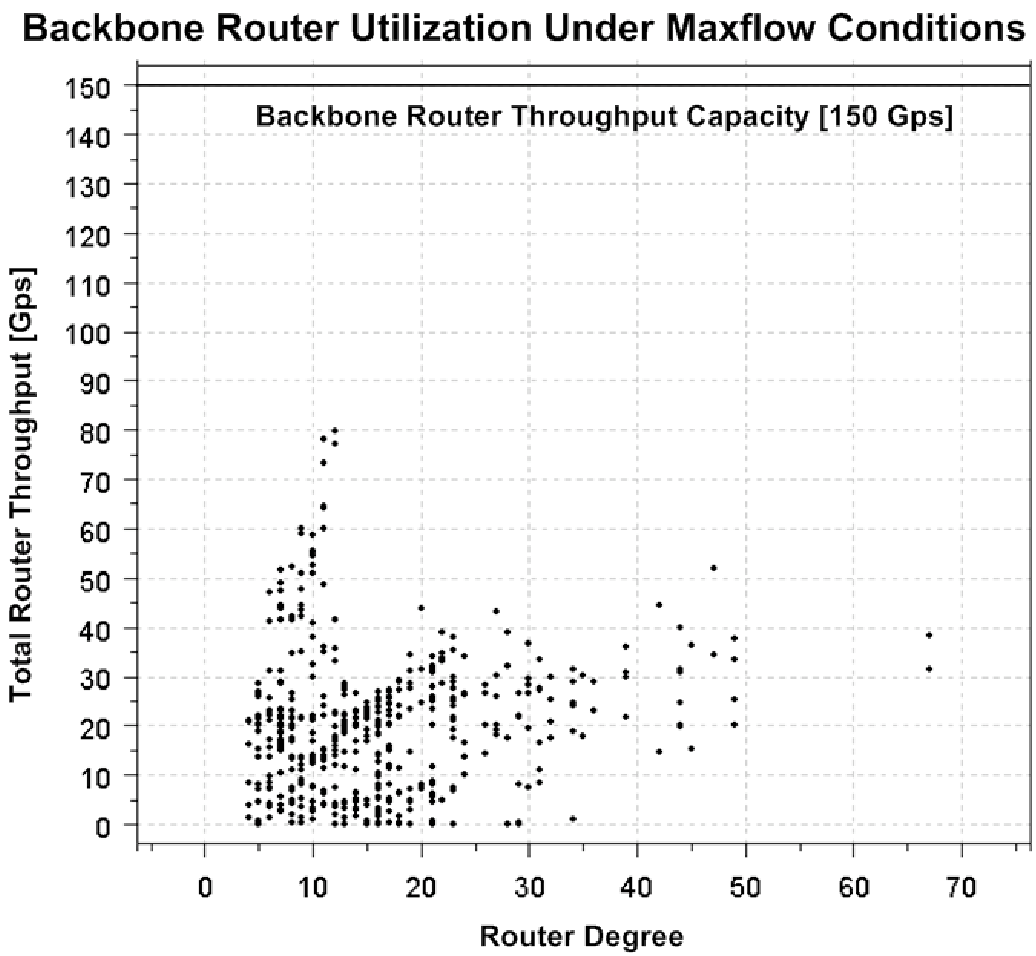

Figure 32. Individual Backbone Router Utilization Under Maxflow Conditions. The backbone routers include routers from all eight MSA subsets and each backbone generation model. Backbone router utilization depends upon the topology structure and traffic engineering used in the network. The wide variation in utilization with a majority of routers being used indicates reasonable resource allocation. Nearly all (99.65%) of backbone routers are utilized, some more then others, with the vast majority under 50% utilization.

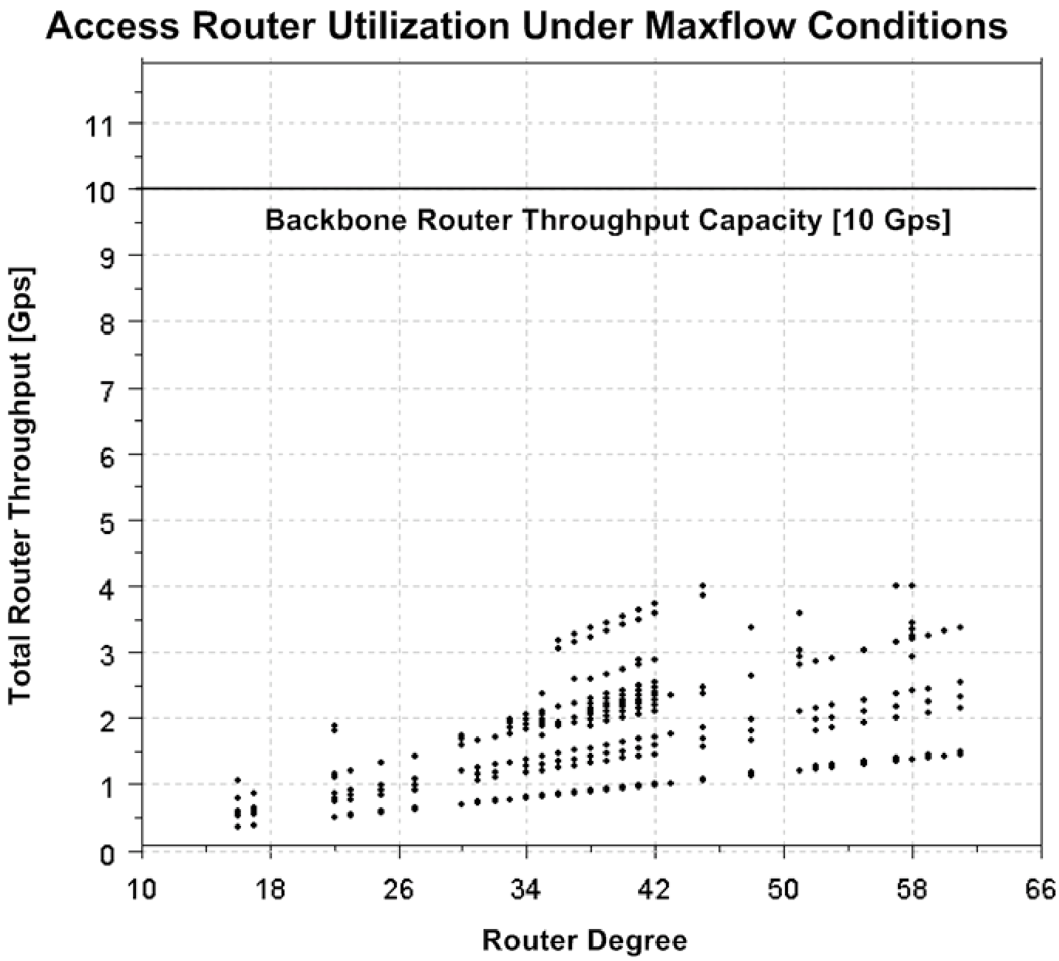

Figure 33. Individual Access Router Utilization Under Maxflow Conditions. Because access routers demand traffic in proportion to the number of customers, the total utilization of access routers increases linearly with customer count. Routers from each subset and backbone model (heuristic, equal cost or equal flow) lie on separate lines.

THIS PAGE INTENTIONALLY LEFT BLANK

V. CONCLUSIONS

In this thesis, we have reverse-engineered network design principles from real world ISP topology and census population data. We have used these design principles to build a topology generation methodology and supporting models. We then used this topology generation process to produce realistic router-level topology maps of different sizes ranging from small regional maps to large national networks. Finally, we evaluated these topologies for cost and throughput performance to (1) validate that the generated topologies are in fact realistic and consistent with what we know about Internet networks and (2) to compare and contrast heuristic and optimal model solutions.

The network topology process and models presented in this thesis do produce realistic models that reflect the observed structure of real ISP topologies. We validate this primarily by throughput analysis and measuring the delivered bandwidth to each customer in the network.

We found that, at the backbone level of representation, optimal design models were able to improve upon as least one of the performance objectives, cost or throughput, by fixing the other. At the router-level representation, cost or throughput improvement did not always correspond to the backbone representation results. For example, the equal cost model throughput was higher then the heuristic and the equal flow cost equal to the heuristic, at the backbone representation level. For the same backbone solutions, represented at the router-level, the heuristic solution might have higher throughput then the equal cost and the equal flow solutions as in subset 3. More work is required to understand why this is so.

In addition to the numerical results, we have developed in this thesis a decision support tool using EXCEL/VBA and GAMS/CPLEX. This tool provides a computational environment where researchers can continue to explore the relationships between network topology and network functionality.

The work in this thesis is based upon several assumptions. We have assumed that the design motifs observed in AS 7018 represent good engineering practice and reflect the structure of the Internet broadly. While we believe the former to be true, we know the latter is not. There are other "styles of design" that may result in dramatically different topologies. For example, anecdotal and empirical evidence for AS 1239 (Sprintlink) suggests that the backbone follows a ring-based design (as opposed to hub-and-spoke) and the internal POP structure follows a hypercube (instead of hierarchical) design. The methodology presented in this thesis would work equally well to incorporate those alternate design motifs, but additional modeling work would be required to include these options.

We have assumed that routers are either one of two types, backbone or access. We know this is not true and many additional types of routers exist, even in AS 7018. For example, terabit backbone router (TBR) pairs are found in several of the larger POPs in AS 7018. In addition, we recognize that the cost and capacity values used as input to our models do not reflect actual equipment, but we have made every attempt to ensure that they are both *externally* consistent (approximate to real equipment, as in Alderson et al. 2004) and *internally* consistent (relative to other parameter values in our model). We have tried to apply this approach of "realistic but fictitious" modeling throughout.

We have assumed a gravity flow model of network traffic in which each pair of communicating routers exchange traffic in proportion to the product of their customer connections. This was also generalized to the backbone topology design were each pair of nodes communicated in proportion to the product of the number of customers in the nodes. In reality, proportionate flow between all customers on the network is not accurate but suffices for large-scale capacity analysis.

Relaxation of any of these assumptions provides many opportunities for future work.

APPENDIX

Full Data Set

Index	MSA	MSA Name	Population	Lat	Lon
1	10420	Akron, OH	694,960	41.1	-81.5
2	10580	Albany-Schenectady-Troy, NY	825,875	42.7	-73.8
3	10740	Albuquerque, NM	729,649	35.1	-106.7
4	12060	Atlanta-Sandy Springs-Marietta, GA	4,247,981	33.7	-84.4
5	12420	Austin-Round Rock, TX	1,249,763	30.3	-97.7
6	12580	Baltimore-Towson, MD	2,552,994	39.3	-76.6
7	13644	Bethesda, MD	1,068,618	39.0	-77.0
8	13820	Birmingham-Hoover, AL	1,052,238	33.5	-86.8
9	14460	Boston, MA	4,391,344	42.4	-71.1
10	14860	Bridgeport-Stamford-Norwalk, CT	882,567	41.2	-73.2
11	15380	Buffalo-Niagara Falls, NY	1,170,111	42.9	-78.9
12	15804	Camden, NJ	1,186,999	39.9	-75.1
13	16580	Champaign-Urbana, IL	210,275	40.1	-88.2
14	16740	Charlotte-Gastonia-Concord, NC	1,330,448	35.2	-80.8
15	16974	Chicago, IL	7,628,412	41.9	-87.7
16	17140	Cincinnati-Middletown, OH	2,009,632	39.2	-84.5
17	17460	Cleveland-Elyria-Mentor, OH	2,148,143	41.5	-81.7
18	17820	Colorado Springs, CO	537,484	38.8	-104.8
19	17900	Columbia, SC	647,158	34.0	-81.0
20	19124	Dallas, TX	3,451,226	32.8	-96.8
21	19340	Davenport-Moline-Rock Island, IA	376,019	41.5	-90.6
22	19380	Dayton, OH	848,153	39.8	-84.2
23	19740	Denver-Aurora, CO	2,157,756	39.7	-105.0
24	19780	Des Moines-West Des Moines, IA	481,394	41.6	-93.6
25	19804	Detroit, MI	2,061,162	42.3	-83.0
26	20764	Edison, NJ	2,173,869	40.3	-74.3
27	22744	Fort Lauderdale, FL	1,623,018	26.1	-80.1
28	23104	Fortworth, TX	1,710,318	32.7	-97.3
29	24340	Grand Rapids-Wyoming, MI	740,482	43.0	-85.7
30	24660	Greensboro-High Point, NC	643,430	36.1	-79.8
31	25420	Harrisburg-Carlisle, PA	509,074	40.3	-76.9
32	25540	Hartford-West Hartford-East Hartford, CT	1,148,618	41.8	-72.7
33	26180	Honolulu, HI	876,156	21.3	-157.9
34	26420	Houston-Sugar Land-Baytown, TX	4,715,407	29.8	-95.4
35	26900	Indianapolis-Carmel, IN	1,525,104	39.8	-86.2
36	27260	Jacksonville, FL	1,122,750	30.3	-81.7
37	28140	Kansas City, MO	1,836,038	39.1	-94.6
38	28700	Kingsport-Bristol-Bristol, TN	298,484	36.7	-82.0
39	29820	Las Vegas-Paradise, NV	1,375,765	36.2	-115.1
40	30780	Little Rock-North Little Rock-Conway, AR	610,518	34.7	-92.3
41	31084	Los Angeles, CA	9,519,338	33.9	-118.3
42	31140	Louisville/Jefferson County, KY	1,161,975	38.3	-85.8
43	31340	Lynchburg, VA	228,616	37.4	-79.1
44	31540	Madison, WI	501,774	43.1	-89.4
45	31700	Manchester-Nashua, NH	380,841	43.0	-71.5
46	32820	Memphis, TN	1,205,204	35.1	-90.0
47	33124	Mimai, FL	2,253,362	25.8	-80.2
48	33340	Milwaukee-Waukesha-West Allis, WI	1,500,741	43.0	-87.9
49	33460	Minneapolis-St. Paul-Bloomington, MN	2,968,806	45.0	-93.3
50	34980	Nashville-Davidson--Murfreesboro--Franklin, TN	1,311,789	36.2	-86.8

Index	MSA	MSA Name	Population	Lat	Lon
51	35004	Long Island, NY	2,753,913	40.8	-73.1
52	35084	Newark, NJ	2,098,843	40.8	-74.4
53	35380	New Orleans-Metairie-Kenner, LA	1,316,510	30.0	-90.1
54	35644	New York, NY	11,296,377	40.7	-74.0
55	36084	Oakland, CA	2,392,557	37.8	-122.3
56	36420	Oklahoma City, OK	1,095,421	35.5	-97.5
57	36540	Omaha-Council Bluffs, NE	767,041	41.3	-95.9
58	36740	Orlando-Kissimmee, FL	1,644,561	28.5	-81.4
59	37964	Philadelphia, PA	3,849,647	40.0	-75.2
60	38060	Phoenix-Mesa-Scottsdale, AZ	3,251,876	33.4	-112.1
61	38300	Pittsburgh, PA	2,431,087	40.4	-80.0
62	38860	Portland-South Portland-Biddeford, ME	487,568	43.7	-70.3
63	38900	Portland-Vancouver-Beaverton, OR	1,927,881	45.5	-122.7
64	39300	Providence-New Bedford-Fall River, RI	1,582,997	41.8	-71.4
65	39580	Raleigh-Cary, NC	797,071	35.8	-78.6
66	40060	Richmond, VA	1,096,957	37.6	-77.5
67	40140	Riverside-San Bernardino-Ontario, CA	3,254,821	34.0	-117.4
68	40380	Rochester, NY	1,037,831	43.2	-77.6
69	41180	St. Louis, MO	2,721,491	38.7	-90.4
70	41620	Salt Lake City, UT	968,858	40.8	-111.9
71	41700	San Antonio, TX	1,711,703	29.4	-98.5
72	41740	San Diego-Carlsbad-San Marcos, CA	2,813,833	32.7	-117.2
73	41860	San Francisco, CA	4,123,740	37.8	-122.4
74	41940	San Jose-Sunnyvale-Santa Clara, CA	1,735,819	37.4	-122.1
75	42044	Anaheim, CA	2,846,289	33.8	-117.9
76	42644	Seattle, WA	2,343,058	47.6	-122.3
77	43780	South Bend-Mishawaka, IN	316,663	41.7	-86.3
78	44060	Spokane, WA	417,939	47.7	-117.4
79	44180	Springfield, MO	368,374	37.2	-93.3
80	45060	Syracuse, NY	650,154	43.0	-76.1
81	45300	Tampa-St. Petersburg-Clearwater, FL	2,395,997	27.9	-82.5
82	45940	Trenton-Ewing, NJ	350,761	40.2	-74.7
83	46060	Tucson, AZ	843,746	32.2	-110.9
84	46140	Tulsa, OK	859,532	36.2	-96.0
85	47260	Virginia Beach-Norfolk-Newport News, VA	1,576,370	36.8	-76.3
86	47644	Warren, MI	2,391,395	42.5	-83.2
87	47894	Washington D.C.	3,727,565	38.9	-77.1
88	48424	West Palm Beach, FL	1,131,184	26.7	-80.1
89	49340	Worcester, MA	750,963	42.3	-71.8

Data Subsets

Subset	Included MSAs	Description
1	1-5, 13, 37	Small
2	39*, 41, 55, 60, 67, 72-75, 83	Southern California
3	4, 8, 12, 15-17, 19, 31, 38, 49, 51, 52, 54, 85	Chicago-Atlanta-New York
4	3, 18, 23, 39, 41, 55, 60, 63, 67, 70, 72-76, 78, 83	West
5	1-2, 6-17, 19, 21-22, 24-27, 32, 35-36, 38, 43-45, 47-49, 51-54, 57-59, 61-62, 64-66, 68, 77, 80, 82, 85-89	East
6	1-4, 6-8, 10-32, 35-36, 38-57, 61-71, 73-86, 88-89	Spoke Heavy
7	4-6, 9, 15-17, 20, 23, 25-28, 34-35, 37, 39, 41, 47, 49, 51-52, 54-55, 58-61, 63-64, 67, 69, 71-76, 81, 85-87	Hub Heavy
8	1-89	All MSAs
*Node 39 in subset 2 has a weight of 2.0		

LIST OF REFERENCES

Alderson, D., L. Li, W. Willinger, and J. Doyle. 2005, “Understanding Internet Topology: Principle, Models, and Validation,” *IEEE/ACM Transactions on Networking*, Volume 13, Number 6, pp. 1205-1218.

Alderson, D. 2008, “Catching the ‘Network Science’ Bug: Insight and Opportunity for the Operations Researcher,’ *Operations Research* (in press).

Alderson, L. Li, D., J. Doyle, and W. Willinger. 2004, “A First-Principles Approach to Understanding the Internet’s Router Level Topology,” *Proceedings of the 2004 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pp. 3–14. New York: ACM Press.

Anderson, D., J. Doyle, R. Govindan, and W. Willinger. 2003, "Toward an Optimization-Driven Framework for Designing and Generating Realistic Internet Topologies," *ACM SIGCOMM Computer Communications Review* Volume 33 Number 1, pp. 41-46.

Barkley, T., 2008. "An Attacker-Defender Model for IP-Based Networks." Master’s Thesis in Operations Research, Naval Postgraduate School, Monterey, CA, March 2008.

Brown, G., M. Carlyle, J. Salmeron, K. Wood. 2006, “Defending Critical Infrastructure.” *Interfaces*, Volume 36, Number 6, pp. 530-544.

Cisco Catalog. 2003. Master Contract for the State of Washington. http://techmall.dis.wa.gov/master contracts/intranet/routers switches.asp (last accessed June 2008).

Spring, N., R. Mahajan, D. Wetherall, and T. Anderson. 2004, “Measuring ISP Topologies with Rocketfuel,” *IEEE/ACM Transactions on Networking* 12, 2–16.

U.S. Census Bureau, 2007, Current List of Metropolitan Statistical Areas, http://www.census.gov/population/www/estimates/metrodef.html (last accessed June 2008).

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
   Ft. Belvoir, Virginia

2. Dudley Knox Library
   Naval Postgraduate School
   Monterey, California

3. Marine Corps Representative
   Naval Postgraduate School
   Monterey, California

4. Director, Training and Education, MCCDC, Code C46
   Quantico, Virginia

5. Director, Marine Corps Research Center, MCCDC, Code C40RC
   Quantico, Virginia

6. Marine Corps Tactical Systems Support Activity (Attn: Operations Officer)
   Camp Pendleton, California

7. Director, Studies and Analysis Division, MCCDC, Code C45
   Quantico, Virginia

8. Professor David L. Alderson
   Naval Postgraduate School
   Monterey, California

9. Professor W. Matthew Carlyle
   Naval Postgraduate School
   Monterey, California